

Windows 8.1 Preview

Windows IT Pro

A PENTON PUBLICATION

AUGUST 2013 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Power Up Your File Servers with SMB 3.0

DHCP Fault Tolerance
in Windows Server 2012

Explore PowerShell Variables

IPv6 in Windows 8 and
Windows Server 2012

Michael Otey
on Hyper-V Server 2012



1&1 DYNAMIC CLOUD SERVER PRICE CONTROL



AS LOW AS

\$0-06
PER HOUR*
PLUS \$60 OFF
FIRST MONTH



COMPLETE COST CONTROL

- **Full transparency** with accurate hourly billing
- **Parallels® Plesk Panel 11 included**, with unlimited domains



FULL ROOT ACCESS

- The complete functionality of a root server with dedicated resources



MAXIMUM FLEXIBILITY

- Configure the Processor Cores, RAM and Hard Disk Space
- Add up to 99 virtual machines



FAIL-SAFE SECURITY

- Redundant storage and mirrored processing units reliably protect your server



Parallels
Plesk Panel



DOMAINS | E-MAIL | WEB HOSTING | eCOMMERCE | SERVERS

Call 1 (877) 461-2631 or buy online

1and1.com

* Other terms and conditions may apply. Visit www.1and1.com for full promotional offer details. Customer is billed monthly for minimum configuration (\$0.06/ hour * 720 hours = \$43.20/ month minimum). A \$60 credit (valid only for the first month) will be applied to your first month of service. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet, all other trademarks are the property of their respective owners. © 2013 1&1 Internet. All rights reserved.



Windows IT Pro Store

eLearning Classes

eBooks

On-Demand Training

Project Plans

Posters

Videos

Plus you can **RENEW** your subscription or
UPGRADE to VIP membership while
you're there!

Stop by the store today!

Windows[®]ITPro

COVER STORY ▼

Enabling the Next-Generation File Server with SMB 3.0

— John Savill

39

For Windows Server 2012, Microsoft wanted to make Server Message Block (SMB) a file-level storage protocol that could be used for crucial enterprise workloads such as Microsoft Hyper-V and SQL Server. To make this shift, some major changes to the SMB protocol were required.

Features

50 Windows Server 2012:
Making DHCP Highly Available
Orin Thomas

69 PowerShell Basics: Variables
Bill Stewart

77 IPv6 Support in Windows 8
and Windows Server 2012
John Howie

Products

95 New & Improved

Interact

89 Ask the Experts

In Every Issue

99 Advertiser Directory

99 Directory of Services

99 Vendor Directory

Chat with Us



Facebook



Twitter



LinkedIn

Columns



6

[Need to Know](#)

Windows 8.1 Preview

Paul Thurrott



13

[Windows Power Tools](#)

Simplifying Services with Managed Service Accounts

Mark Minasi



16

[Top 10](#)

Important Features in Microsoft's Free Hyper-V Server 2012

Michael Otey



19

[Enterprise Identity](#)

What Are OAuth 2.0 and OpenID Connect?

Sean Deuby



23

[What Would Microsoft Support Do?](#)

Troubleshooting Windows Server 2012 Failover Clusters

John Marlin

Editorial

Editorial Director: Megan Keller
Editor-in-Chief: Amy Eisenberg
Senior Technical Director: Michael Otey
Technical Director: Sean Deuby
Senior Technical Analyst: Paul Thurrott
IT Community Manager: Rod Trent
Systems Management, Networking,
Hardware: Jason Bovberg
Scripting: Blair Greenwood
SharePoint, Active Directory, Security,
Virtualization: Caroline Marwitz
SQL Server, Developer Content:
Megan Keller
Managing Editor: Lavon Peters
Editorial SEO Specialist: Jayleen Heft

Senior Contributing Editors

David Chernicoff, Mark Minasi,
Tony Redmond, Paul Robichaux,
Mark Russinovich, John Savill

Contributing Editors

Alex K. Angelopoulos, Michael Dragone,
Jeff Felling, Brett Hill, Dan Holme,
Darren Mar-Elia, Eric B. Rux,
William Sheldon, Curt Spanburgh,
Bill Stewart, Orin Thomas, Douglas Toombs,
Ethan Wilansky

Art & Production

Senior Graphic Designer: Matt Wiebe
Director of Production: Dylan Goodwin
Group Production Manager:
Julie Jantzer-Ward
Project Manager: Adriane Wineinger
Graphic Specialist: Karly Prickett

Advertising Sales

Technology Market Leader: Peg Miller
Key Account Director:
Chrissy Ferraro • 970-203-2883
Account Executives:
Megan Key • 970-203-2844
Barbara Ritter • 858-367-8058
Cass Schulz • 858-357-7649

Client Services

Senior Client Services Manager:
Michelle Andrews • 970-613-4964
Ad Production Coordinator: Kara Walby

Marketing & Circulation

Customer Service • 800-793-5697
Senior Director, Marketing Analytics:
Tricia Syed

Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

Corporate

Chief Executive Officer:
David Kieselstein
Chief Financial Officer/Executive Vice
President: Nicola Allais



List Rentals

MeritDirect
333 Westchester Avenue,
White Plains, NY 10604

Reprints

Reprint Sales:
Wright's Media • 877-652-5295

Windows IT Pro, August 2013, Issue No. 228,
ISSN 1552-3136. *Windows IT Pro* is published monthly by
Penton. Copyright ©2013 Penton. All rights reserved. No
part of this publication may be reproduced or distributed
in any way without the written consent of Penton.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525,
800-621-1544 or 970-663-4700. Customer Service:
800-793-5697.

We welcome your comments and suggestions about the
content of *Windows IT Pro*. We reserve the right to edit all
submissions. Letters should include your name and
address. Please direct all letters to letters@windowsitpro.com. IT pros interested in writing for *Windows IT Pro* can
submit articles to articles@windowsitpro.com.

Program Code: Unless otherwise noted, all programming
code in this issue is ©2013, Penton, all rights reserved.
These programs may not be reproduced or distributed
in any form without permission in writing from the
publisher. It is the reader's responsibility to ensure
procedures and techniques used from this publication are
accurate and appropriate for the user's installation. No
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®
are trademarks or registered trademarks of Microsoft
Corporation in the United States and/or other countries
and are used by Penton, under license from owner.
Windows IT Pro is an independent publication not
affiliated with Microsoft Corporation. Microsoft
Corporation is not responsible in any way for the editorial
policy or other contents of the publication.

Windows IT Pro

Windows 8.1 Preview



Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for *Windows IT Pro UPDATE*, and a daily Windows news and information newsletter called *WinInfo Daily UPDATE*.

Email



Twitter



Website



By the time you read this, Microsoft will have delivered its first and only prerelease version of the Windows 8.1 update for both [Windows 8](#) and Windows RT. Dubbed the Windows 8.1 Milestone Preview (MP) internally, this release is meant to encapsulate the Developer Preview, Consumer Preview, and Release Preview milestones of previous Windows versions in a single release, a sign of this product's much quicker development schedule. It offers a lot, both for users and for businesses that, until now, were ready to give Windows 8 and Windows RT a pass.

The Should-Have-Been Original Release?

After spending two weeks with a near-final version of the Windows 8.1 Preview on several PCs, I've found that this update does indeed smooth over many of the rough patches in Windows 8, providing a more cohesive and complete experience than what Microsoft first delivered late last year. Is this the release that Microsoft should have shipped originally? It's perhaps trite to suggest such a thing. But yes, that's obviously the case.

Windows 8.1 represents a lot of things for Microsoft. Its name suggests that it's a new Windows version and not just an update, and perhaps not coincidentally even Microsoft has sometimes referred to it as such. You might draw some not-unreasonable comparisons between this release and Windows 3.1 from 20 years ago; like its ancient predecessor, this release is essentially a fine-tuning of an earlier major update that was perhaps released a bit too quickly.

Windows 8.1 is also our first peek at what the Windows team can accomplish in a year. Based on conversations with numerous sources throughout Microsoft, it's become fairly obvious that this team has moved uncomfortably into the online services model that has been embraced by other parts of the company, such as Office and Windows

Server. (The phrase “kicking and screaming” was used in one humorously descriptive note.)

And while there are ongoing murmurs of discontent about being forced to update Windows so quickly, this much is clear: If ever a product needed this rapid an update, it's Windows 8.

Of course, Windows 8.1 can't change the fundamental disconnect in the underlying product, which remains a strange amalgam of the classic Windows desktop environment we've been using since Windows 95 and a brand-new touch-first mobile environment that I will continue calling Metro. Jammed together like a technological Frankenstein's monster, these two environments established an uneasy coexistence in the original release of Windows 8 (and RT), with jarring transitions between the two that were made all the more painful by the fact that most users were still using PCs without touch capabilities.

Finally, Fixing Windows 8

Looked at from a high level, some of the biggest changes in Windows 8.1 are designed to address these problems. For example, Microsoft designed the system so that those who wanted to stay in the desktop environment—typically those with classic PC hardware—could skip the Metro Start screen at boot and go directly to the desktop. And while in the desktop, these users can mostly ignore (or even disable) annoying Metro-style UI elements.

There's even a modern take on the Start button (which was removed in Windows 8 and RT), with the option for it to trigger an ancillary All Apps view rather than the reviled Start screen, if desired. (Sad trombone moment: Like the Start screen, All Apps is a Metro experience, too.)

On the flip side, those who actually want to stay in the Metro environment—yes, I'm told such people exist—can do so more easily now as well. These users, who will most typically be using a new generation of tablets and hybrid PC devices, will find that Windows 8.1

adds far more customization settings to the Metro-based PC Settings interface, negating the need (as in the original release of Windows 8) to find and use the desktop-based Control Panel.

Start Screen Changes

But it's not just about settings. Metro is also a more elegant environment now than it was in that first Windows 8 release. The Start screen sports optional animated backgrounds and can use the same wallpaper as the desktop if you want, cutting down on that annoying jarring effect when you switch between them. The Start screen tiles now support four sizes, instead of just two (see Figure 1), and the larger sizes offer more live information, often attractively presented, and relief for those on high-resolution displays.

Figure 1
Windows 8.1's Start
Screen



Those who want to customize the Start screen no longer need to learn about and locate hidden commands found in far-off places. You won't inadvertently move tiles around, as it was easy to do in the initial release, and you can now customize Start directly from Start, seeing your changes as you apply them. And many more changes abound, including a literal rainbow of color choices, instead of the stock few in the first release.

Eventually, of course, you'll need to get past the Start screen. And again looking at this release from a high level, I'd choose its new and deeper integration with SkyDrive as perhaps the biggest and most important change.

If you're familiar with Windows 8, you know that the system lets you sign in with a Microsoft account, then sync a limited set of settings across all of your PCs through SkyDrive. That is literally the extent of SkyDrive integration in Windows 8, and if you want more—such as PC-SkyDrive file sync—you need to find and install an application. (And that application isn't available on RT.)

SkyDrive Integration

In Windows 8.1, SkyDrive becomes a profound and truly integrated part of the Windows experience. During setup, you're asked if you want to integrate your SkyDrive storage with Windows. If you choose to do so, your SkyDrive storage is integrated into the file system, and if you navigate through that folder structure (in a SkyDrive folder in your user profile), you will believe that all of your SkyDrive-based files and folders have been downloaded and synced to the PC. But that's not the case.

Instead, what you're seeing is a set of new shortcut types that look and work like the actual files. If you're online, you can simply open them as usual: Microsoft Word documents open in Word, Adobe Photoshop documents open in Photoshop. Everything works.

Going offline? Simply right-click any file or folder and choose the new *Make available offline* command from the menu that appears. It's powerful, granular, and speedy.

Back in Metro, SkyDrive works much like it does on Windows Phone. You can configure the system to automatically back up all pictures and videos to SkyDrive that you take with the device's internal camera, and at full resolution, if desired. There are far more settings being synced between your PCs and devices, an enhancement of the functionality that debuted with Windows 8.

Metro App Changes

Microsoft updated several of the built-in Metro-style apps in Windows 8 and RT earlier this year, and Windows 8.1 will arrive with some more changes. In the Preview build, we see a dramatically updated Photos app, for example, though the early version I've used has lost some useful functionality (only temporarily, I hope) around photo acquisition and online services integration. Xbox Music gets a prettier and more usable UI that more closely resembles traditional jukebox software such as iTunes and uses a single screen instead of Windows 8's rambling panoramic experiences.

There are new apps, too. Bolstering the selection of surprisingly useful and beautiful Bing apps from the initial release—Bing, News, Sports, Finance, Travel, and Maps—Windows 8.1 includes two new Bing apps: Food & Drink and Health & Fitness.

There are also new utility apps such as Calculator, Help & Tips, Reading List, Scan, and Sound Recorder. Reading List is a news reader-type app, which looks attractive. And yes, Scan is exactly what it sounds like: A Metro-based scanner utility.

Desktop Changes

Desktop users don't get as many updates as do Metro users, but let's be fair: Most of the big issues were on the Metro side. But in addition to the ability to boot to the desktop and show All Apps instead of Start, there are other improvements for us Luddites.

Windows 8.1 lets you sort that All Apps view so that desktop applications are listed first, before Metro apps. You can disable the Metro-based Charms and Switcher interfaces. And the secret power-user menu (WinKey + X) now has more commands.

But the biggest desktop change is the introduction of new display scaling capabilities. Available via Display in Control Panel, this new capability automatically scales the desktop according to the screen resolution and physical size of the display, making, say, small screen or high-resolution devices such as Surface Pro instantly more usable.

This capability has other niceties: You can manually override the automatic settings, and, although it doesn't appear to work in the Preview release, there's an option for maintaining different display scaling settings for each display attached to your PC. So using Surface Pro as an example again, you might see (or configure) 150 percent scaling on the device's tiny internal display but 100 percent scaling on an attached 27-inch display. Can I get a Hallelujah?

Business Features

Amazingly, and inexplicably, Windows 8.1 also arrives with a surprising range of new features aimed directly at businesses. Querying Microsoft about this—surely most businesses are planning to skip this product generation, I said—I was told that the firm really does expect businesses to roll out Windows 8.x alongside Windows 7, using the former on devices instead of traditional PCs. We'll see whether that view translates into reality, but there's no denying the effort.

Windows 8.1 adds such networking features as NFC tap-to-pair printing, Wi-Fi Direct printing, Miracast wireless display, broadband tethering (Internet sharing), and massive improvements to the built-in VPN capabilities, including compatibility with several third-party VPNs (sadly, not Cisco). A new version of Internet Explorer, IE11, provides “faster page load times, side-by-side browsing of your sites, 3D graphics, enhanced pinned site notifications, reading view and app settings like favorites, tabs and settings sync across all your Windows 8.1 PCs.”

Windows 8.1-based devices (including those using RT) will support selective remote wipe, so that users who bring their own machines to work and decide to remove them from corporate control won't lose their personal data. Windows 8.1 features a new WorkPlace Join capability so that users can easily join a domain from Metro and set up their devices for policy-based corporate control.

A new Work Folders feature (part of [Windows Server 2012 R2](#)) will provide SkyDrive-style client sync with corporate document libraries.

And admins can now control far more of what users can see and do on their Windows 8.1-based devices, including what apps appear and how the Start screen is laid out.

Is It Enough?

While it's interesting to see how much the Windows team can accomplish in just a year, questions remain. Windows 8 got off to a rough start and Windows RT might be charitably described as a disaster so far. The changes in Windows 8.1 are sometimes major—especially for Windows RT usage at work—but as is the case with many Windows updates, it's the combination of many minor changes that puts this release over the top.

You should see for yourself. The Windows 8.1 Preview is freely available from the Microsoft website and will update any Windows 8 PC or Windows RT device. Note, however, that those who do install the Preview will lose any installed desktop applications or Metro-style apps when they later upgrade to the final release, which at the time of writing was scheduled to be completed in August. ■

Simplifying Services with Managed Service Accounts

MSAs will be the reason you decide to finally learn PowerShell

Sometimes when I talk about Active Directory (AD) administration with PowerShell, people seem open to the idea of learning PowerShell, but I know they're privately hoping they'll never really need to. They might say, "I don't really need to learn this stuff to use AD in [Windows Server 2012](#) or Windows Server 2008 R2, do I?" I have to answer truthfully that, yes, you can still do a lot of work in the GUI—I'm just not quite sure why you'd want to! Just as a quick example, setting my title to instructor is a lot quicker to accomplish with

```
set-aduser mark -title 'teacher'
```

than using the Active Directory Users and Computers snap-in or the Active Directory Administrative Center. And unlocking Larry's account is a heck of a lot easier with

```
unlock-adaccount larry
```

than by spelunking through Active Directory Users and Computers. Directing any query to a Global Catalog server is far simpler with PowerShell (just add `-server servername:3268` to any `get-aduser` query) than with the GUI. Sometimes these kinds of examples will convince the skeptics, but not always. In those cases, I've got to pull out the big guns: managed service accounts (MSAs).



Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.



Email



Twitter



Website

Have you ever had to create a domain account because you needed to run a service on a server, and it needed its own account to run under (or perhaps one to run an IIS application pool under)? If so, the following story might sound familiar. You're in the middle of a 25-step setup process for an important server-based application, and it wants you to create a domain account. Or perhaps its Setup program informs you that it has created an account for the new service to run under. You either create the account or let the Setup program create the account, then you move along to the next step and eventually get the new service running. All is now well, and everyone loves the new service—for a while.

One day, you walk into the office, and everyone's hair is on fire. The service you set up six weeks earlier is no longer working, no one can figure it out, and it dawns on you that your domain password policy requires a new password every six weeks. Eek! So you reset the password, look around to ensure that no security people are watching, and select the *Password never expires* check box.

Alternatively, if you have Server 2012 or Server 2008 R2, you could skip creating a domain user account to run the service under and instead set up an MSA. As long as you have at least one Server 2008 R2 domain controller (DC) and you're running that service on a Server 2012 or Server 2008 R2 member server, you simply create an MSA and configure the service to run under the MSA (leave the password field in the Services snap-in blank, and it'll get filled in automatically) and AD. Then, the MSA account and the member server will create a new password once a month—with no need for human intervention.

I have no idea why most Server 2012/Server 2008 R2 admins have never heard of MSAs, but when I tell the preceding story to my skeptical AD learners, their eyes get big. "What does that have to do with PowerShell?" they ask. "Everything!" I answer, because for some reason the only way you can create an MSA is with PowerShell.

The PowerShell "noun" that describes MSAs is `ADServiceAccount`, and if you've spent a couple of hours with PowerShell, you'll probably

guess that you can create an MSA with *new-adserviceaccount*. For example, to create an account named *svc1*, you would type

```
new-adserviceaccount -name svc1
```

and if you're doing this on Server 2012, you'd add *-RestrictToSingleComputer*, as in

```
new-adserviceaccount -name svc1 -RestrictToSingleComputer
```

Then, you would walk over to the member server where you'll be running the service or do an *Enter-PSSession* to that system (the *-computername* parameter doesn't work on this cmdlet) and essentially "introduce" the managed service account to the member server with the *install-adserviceaccount* cmdlet, followed by a space and the name of the managed service account, as in

```
install-adserviceaccount svc1
```

At that point, you'd need only tell the service to run under the account name. To do so, as I mentioned earlier, simply open the Services snap-in and—where you'd fill in the account that the service runs under—just fill in the MSA account name, and don't put anything in the password area. An MSA behaves sort of like a machine account, so add a dollar sign (\$) to the end of the name. For example, in the case of our *svc1* account, if it were in a domain with a NetBIOS name of *bigfirm*, it would be

```
bigfirm\svc1$
```

The rest is automatic. Now, I can't swear this is true, but I think MSAs might have been the "closer" for a few folks when it comes to learning PowerShell. ■

Important Features in Microsoft's Free Hyper-V Server 2012

Take advantage of enhancements to Hyper-V virtualization support



**Michael
Otey**

is senior technical director for
Windows IT Pro and *SQL
Server Pro*.

Email



In [Windows Server 2012](#), Microsoft significantly enhanced the Hyper-V virtualization support. Although not everyone realizes it, Microsoft also offers a completely free version of its hypervisor, called Hyper-V Server 2012. Let's take a look at the 10 most important features in [Microsoft's free Hyper-V Server 2012](#).

① Guest OS Licenses

One of the most important differences to note between Hyper-V Server 2012 and [Windows Server 2012 Hyper-V](#) is the fact that Hyper-V Server 2012 doesn't include any guest OS licenses. (For more information about Windows Server 2012 licensing, see the [Microsoft Windows Server 2012 Volume Licensing Buyer's Guide](#).) Hyper-V Server 2012 is a great option when you want to run Linux or for virtual desktop infrastructure (VDI) scenarios in which the licensing advantages of the Windows Server 2012 Standard and Datacenter editions don't apply.

② Host Scalability

Hyper-V Server 2012 provides the same scalability that you can get from Windows Server 2012 Hyper-V. Hyper-V Server 2012 supports 320 logical host processors, 4TB of host RAM, 2,048 virtual CPUs per host, and 1,024 active virtual machines (VMs) per host.

③ Guest Scalability

The free Hyper-V Server 2012 supports the same high levels of guest scalability that are found in Windows Server 2012 Hyper-V. Hyper-V Server 2012 provides support for 64 virtual CPUs per VM, support for guest Non-Uniform Memory Access (NUMA), and support for up to 1TB of RAM per VM.

④ Cluster Support

Hyper-V Server 2012 systems can fully participate in Microsoft Windows failover clusters. There's full support in Hyper-V Server 2012 for up to 64-cluster nodes and up to 8,000 VMs.

⑤ Storage Capabilities

Hyper-V Server 2012 shares most of the new Windows Server 2012 Hyper-V storage enhancements. Hyper-V Server 2012 supports up to four virtual Fibre Channel adapters per VM. Hyper-V Server 2012 also supports the new VHDX format with up to 16TB virtual hard disks. In addition, Hyper-V Server 2012 provides Offload Data Transfer (ODX) support for high-performance SAN data transfers and Windows Server 2012 Storage Spaces.

⑥ Advanced Networking

Microsoft introduced NIC teaming using heterogeneous network adapters in Windows Server 2012. Hyper-V Server 2012 inherits the same NIC teaming ability, which you can configure with the Hyper-V PowerShell cmdlets. Hyper-V Server 2012 also supports single root I/O virtualization (SR-IOV). For high-performance VM networking, Hyper-V Server 2012 supports quality of service (QoS), providing the ability to specify the minimum bandwidth available to a VM or a port.

⑦ Dynamic Memory

Hyper-V Dynamic Memory was first introduced with Windows Server 2008 R2 SP1. Dynamic memory increases the server consolidation

ratios that are possible. Hyper-V Server 2012 fully supports memory overcommit. However, like Windows Server 2012 Hyper-V, the guest OS must support hot-add RAM in order to take advantage of dynamic memory.

Hyper-V Server 2012 provides the same scalability that you can get from Windows Server 2012 Hyper-V.

8 VM Mobility

Microsoft first introduced live migration in Windows Server 2008 R2 and significantly enhanced this feature in Windows Server 2012. Hyper-V Server 2012 supports all of the live migration options that are provided in the full Windows Server 2012 Hyper-V implementation. Hyper-V Server 2012 supports Shared Storage Live Migration, Server Message Block (SMB) Live Migration, Shared-Nothing Live Migration, and Storage Live Migration. Multiple concurrent live migrations are supported and all can be run with no end-user downtime. There is also full support for Hyper-V Replica.

9 Hyper-V Extensible Switch

The new Hyper-V Extensible Switch is fully supported in Hyper-V Server 2012. The Hyper-V Extensible Switch supports internal, external, and private switches. There's also support for Private VLANs (PVLANS) and DHCP Guard. Like Windows Server 2012 Hyper-V, Hyper-V Server 2012 is fully extensible and supports multiple filter extensions, capture extensions, and forwarding extensions.

10 Network Virtualization

Microsoft first introduced network virtualization in Windows Server 2012; this feature allows you to extend your networks across different subnets and from your on-premises networks into the cloud. Network virtualization enables you to seamlessly move VMs from on-premises into the cloud and back, with no downtime and no need to change the VM or application's networking. ■

What Are OAuth 2.0 and OpenID Connect?

Two identity frameworks support the next generation of web single sign-on

In “Attention, IT Pros: You Can Help Evolve a Secure Cloud, Too,” I emphasized how important it is for businesses to support open Internet identity standards and to require their vendors to support them. I’m long overdue to bring a couple of relatively new, very important identity frameworks to your attention. Why should you care? These new frameworks—[OAuth 2.0](#) and [OpenID Connect](#)—are the “Kerberos of the cloud.” As with Kerberos, even if you don’t explicitly develop code using them, you must at least know how they work in order to support your users: Those users are utilizing mobile and [cloud](#) apps with increasing frequency.

We all know the [problem with passwords](#). In case you’ve just come back from a dogsled run to the North Pole, all you need to do is ask your family . . . or friends . . . or neighbors . . . or a random stranger on the street. We have far too many user IDs and passwords. We can’t keep track of them, and so we make them very simple so that we can remember them. Then we reuse these simple passwords across many websites in an effort to lessen the confusion. (Be honest: How many of you have perfect password hygiene? If *you* can’t do it, how can you possibly expect Aunt Sally to get it right?) This situation creates the “password anti-pattern,” in which we enter our user ID and password from one site (the identity provider—IdP) to gain access to another site (the service provider—SP). Worse, we’ve become conditioned to this prompt so that we might mistakenly enter it at a bogus malware prompt. In May, Google developer



Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel’s core directory services team. He’s been a directory services MVP since 2004.



Email



Twitter

evangelist [Tim Bray got on his knees and pounded the floor](#), pleading with developers not to require him to create yet another user ID and password for every new website.

This situation is made worse with mobile devices. The difficulties of entering credentials are magnified on the smaller form factor; it's harder to type in passwords on small keyboards (and if you're old enough, without putting on glasses). Thus, it's much more difficult to type in the strong passwords you're supposed to be using. Once inputted, these passwords are often stored insecurely on the device. Then devices get lost. And how many people use PIN codes or some other kind of lock on their devices? A 2011 survey by Confident Technologies found that [fewer than half of mobile device users lock their device](#).

OAuth 2.0

Fortunately, OAuth 2.0 is available to help with your password pain, especially for mobile apps. [An IETF proposed standard](#), OAuth 2.0 is technically an “authorization framework.” But a better description I've heard is that it's an “authorization-centric,” flexible protocol that also supports authentication.

Generally speaking, OAuth 2.0 uses the *Get a token, use a token* methodology, as you see in Figure 1. A user (more specifically, an application

Figure 1
OAuth 2.0 Client
Credentials Flow



on a user's full or mobile client) wants to gain access to a resource protected by the resource server. To authenticate to the resource server, the client must include an access token in its communication to the server. This token is provided by the authorization server. In other words, the client gets a token from the authorization server, then uses the token to authenticate to the resource server, thus gaining access to the resource. You can see how the OAuth 2.0 client credentials flow bears a passing resemblance to the Kerberos credentials flow.

A number of very visible OAuth 2.0 examples are at work today. If you choose to log on to a web app (e.g., [Twitter](#), [TripIt](#)) using your account from an identity provider (e.g., [Google](#)), watch the authentication process carefully and you'll see URLs that contain *oauth*. The OAuth logo itself might even briefly pop up. In his hilariously titled "[Is that a token in your phone in your pocket, or are you just glad to see me?](#)" slide deck, [Brian Campbell](#) explains how OAuth 2.0 works with mobile clients, and [Ping Identity](#) offers a white paper entitled "[The Essentials of OAuth](#)" (registration required) that provides a very clear overview.

I mentioned that OAuth 2.0 is more flexible than Kerberos. It's also more complicated to implement. The *It's an authorization protocol! / It's an authentication protocol!* debate is just one aspect of the confusion. That's where OpenID Connect comes in.

OpenID Connect

OpenID Connect is a simple identity layer on top of OAuth 2.0. It's a specification that organizes how identity providers and relying parties can use OAuth 2.0 to communicate identity data to one another, without having to code a full OAuth implementation. By easing developer pain, the hope (and it seems to have been borne out in practice) is that more developers will use OAuth 2.0 to provide secure authentication. In her article "[OpenID Connect: New, groovy and full of promise](#)," Pamela Dingle provides a good overview of how it works, and Oliver Pfaff delves into the details in his "[OpenID Connect—An Emperor Or Just New \[Clothes\]?](#)" slide deck.

My Standard Is Better Than Your Standard: Open Standards Confusion

Researching this column, I learned just how confusing the open standards world can be when compared with a single vendor's implementation. First, there are many standards organizations, such as the [IETF](#), [OASIS](#), and the [OpenID Foundation](#). Each of these organizations has a variety of proposals, with a committee, working their way through the standards process. Not all protocols are standards, and not all standards are protocols. Further, each proposal, or framework, makes its way through drafts, specifications, proposed standards—and this nomenclature might vary between organizations. One framework from one standards body (e.g., OpenID Connect from the OpenID Foundation) can be designed to work on top of a framework from another standards body (e.g., OAuth 2.0 from the IETF). Naming can be confusing even within an organization; for example, OpenID Connect isn't related to OpenID. And after all this, I'm sure I still got something wrong! ■

Good to Know

In February, Gartner predicted that half of new identities on retail sites will be based on social network identities (e.g., Facebook, Google, Twitter, Microsoft) rather than identities created directly on the retail site. This can only happen with a common method to easily and securely provide identities from IdPs to SPs to authenticate with, and OAuth 2.0 and OpenID Connect seem to be the most popular.

If your daily job doesn't require you to work with external identities, learning about OAuth 2.0 and OpenID Connect is simply a good idea—a bit of knowledge to tuck away for future use. If your company works at all with consumer identities, these two identity frameworks will have some part in your future. Take this introduction and related links, and dig deep enough so that you can judge how they'll impact or empower your job. And if you're confused (as I was) about all these open standards, be sure to check out the sidebar “My Standard Is Better than Your Standard: Open Standards Confusion.” ■

Troubleshooting Windows Server 2012 Failover Clusters

How to get to the root of the problem

In “[Troubleshooting Windows Server 2008 R2 Failover Clusters](#),” I discussed troubleshooting failover clusters—specifically, the locations and tips for where you can go to get the data you need to troubleshoot a problem. Now I’ll discuss some of the improvements made to the troubleshooting tools for [Windows Server 2012](#) failover clusters and show you how to take advantage of those tools.

Introducing the New Event Channels

There are some new event channels for failover clustering to help with troubleshooting. Figure 1 shows all the available channels. Note that the events are specific to the node you’re on.

Knowing the purpose of each event channel can help you find the errors more quickly, which in turn will help you troubleshoot the problem more quickly. Here’s an explanation of each channel:

- FailoverClustering
 - Diagnostic. This is the main log that’s circular in nature and runs anytime the cluster service starts. Events can be read in the Event Viewer if logging is disabled. They can also be converted to text file format.
 - Operational. Any informational cluster events are registered in this log, such as groups moving, going online, or going offline.
 - Performance-CSV. This channel is used to collect information pertaining to the functionality of Cluster Shared Volumes (CSVs).



John Marlin

is a senior support escalation engineer in Windows Commercial Technical Support, focusing on failover clustering. He is a Microsoft Certified Trainer for clustering, delivering to Microsoft and its partners, and is a regular contributor to the [Ask the Core Team](#) blog. He is also a contributor to the new book *Introducing Windows Server 2012* (Microsoft Press).

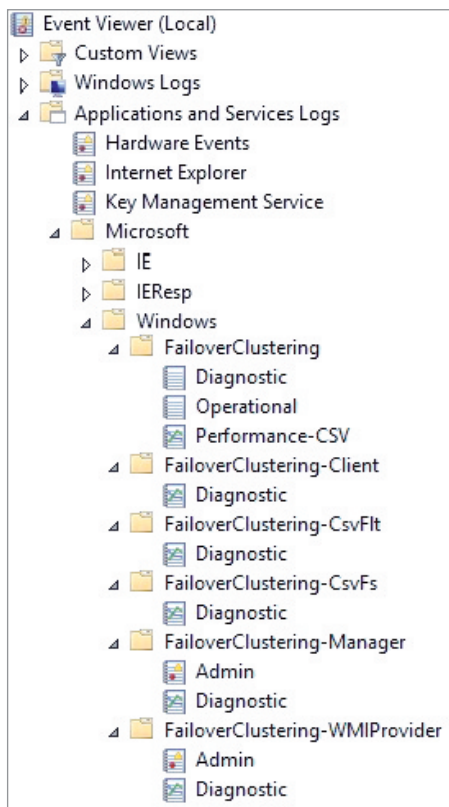


Email



Website

Figure 1
Event Channels for
Failover Clustering in
Server 2012



- FailoverClustering-Client
 - Diagnostic. This channel collects Cluster API trace logging. This log can be useful in troubleshooting the *Create Cluster* and *Add Node Cluster* actions.
- FailoverClustering-CsvFlt (new in Server 2012)
 - Diagnostic. This channel collects trace logging for the CSV Filter Driver (CsvFlt.sys) that is mounted only on the coordinator node for a CSV. This channel provides information regarding metadata operations and redirected I/O operations.
- FailoverClustering-CsvFs (new in Server 2012)
 - Diagnostic. This channel collects trace logging for the CSV File System Driver (CsvFs.sys), which is mounted on all nodes in the cluster. This channel provides information regarding direct I/O operations.
- FailoverClustering-Manager
 - Admin. This channel collects errors associated with dialog boxes and pop-up warnings that are displayed in Failover Cluster Manager.
- FailoverClustering-WMIProvider
 - Admin. This channel collects events associated with the Failover Cluster WMI provider.
 - Diagnostic. This channel collects trace logging associated with the Failover Cluster WMI provider. It can be useful when

troubleshooting Windows Management Instrumentation (WMI) scripts or Microsoft System Center applications.

Using the FailoverClustering-Client/Diagnostic Channel

Because administrators often encounter problems when creating clusters and joining nodes, I want to show you how to use the FailoverClustering-Client/Diagnostic channel. This channel is disabled by default, so it won't be collecting any data. To enable it, you need to right-click the channel and choose Enable Log. The Diagnostic channel will then start collecting data relevant to a join or create operation.

For example, suppose you previously enabled the Diagnostic channel and you're having a problem creating a cluster. To view the data collected, you need to right-click the channel and choose Disable Log. In the FailoverClustering-Client/Diagnostic event log, you see the following events:

Event ID: 2

Level: Error

Description: CreateCluster (1883): Create cluster failed with exception. Error = 8202, msg: Failed to create cluster name CLUSTER on DC \\DC.CONTOSO.COM. Error 8202.

Event ID: 2

Level: Error

Description: CreateClusterNameCOIfNotExists (6879): Failed to create computer object CLUSTER on DC \\DC.CONTOSO.COM with OU ou=Clusters,dc=contoso,dc=com. Error 8202.

Because you have errors, you can use the Net.exe command to see what their status code (8202) means:

```
NET HELPMMSG 8202
```


The command returns the message: *The specified directory service attribute or value does not exist.* With the new features of Server 2012 Failover Clustering, the cluster will be created in the same organizational unit (OU) as the nodes. For the cluster name to be created, the logged-on user must have at least Read and Create Computer Objects permissions. If the user doesn't have those rights, the name won't be created and you'll receive this type of error.

Now suppose you're trying to add a node to the existing cluster and the operation fails. You review the events in the FailoverClustering-Client/Diagnostic log, and see the following:

Event ID: 56

Level: Warning

Description: AsyncNotificationCallback (1463): ApiGetNotify on hNotify=0x00000000021EBCDC0 returns 1 with rpc_error 0

Event ID: 2

Level: Error

Description: SCMStateNotify (837): Repost of NotifyServiceStatusChange failed for node 'NodeX': status = 1168

Although their wording is a bit on the cryptic side, the descriptions give you the information that you need. The description for the first event tells you that a remote procedure call (RPC) error occurred. The description for the second event gives you a status code of 1168. Once again, you can use the Net.exe command to see what that status code means:

```
NET HELPMSG 1168
```

This time, the command returns the message: *Element not found.* When a node tries to join a cluster, the running cluster node needs

to make an RPC connection to the node being added. In this case, it couldn't find the node.

So, from the information returned by the two events, you can deduce that the running cluster node can't make an RPC connection to the node being added because it can't find that node. After further investigation, you discover that the DNS server has an incorrect IP address for the node being added. After you correct the IP address, the node successfully joins the cluster.

Introducing the New Tests in the Validate a Configuration Wizard

Another helpful troubleshooting tool that you can use is the Validate a Configuration Wizard in Failover Cluster Manager. Several new clustering tests have been added in Server 2012. All the new tests for Server 2012 clustering are in bold:

- **Hyper-V (available only if the Hyper-V Role is installed)**
 - **List Hyper-V Virtual Machine Information**
 - **List Information About Servers Running Hyper-V**
 - **Validate Compatibility of Virtual Fibre Channel SANs for Hyper-V**
 - **Validate Firewall Rules for Hyper-V Replica Are Enabled**
 - **Validate Hyper-V Integration Services Version**
 - **Validate Hyper-V Memory Resource Pool Compatibility**
 - **Validate Hyper-V Network Resource Pool and Virtual Switch Compatibility**
 - **Validate Hyper-V Processor Pool Compatibility**
 - **Validate Hyper-V Role Installed**
 - **Validate Hyper-V Storage Resource Pool Compatibility**
 - **Validate Hyper-V Virtual Machine Network Configuration**
 - **Validate Hyper-V Virtual Machine Storage Configuration**
 - **Validate Matching Processor Manufacturers**
 - **Validate Network Listeners Are Running**
 - **Validate Replica Server Settings**

When reviewing
Cluster.log files, it
helps to search for
keywords.

- Cluster Configuration (available only if a cluster is running)
 - List Cluster Core Groups
 - List Cluster Network Information
 - List Cluster Resources
 - List Cluster Volumes
 - List Clustered Roles
 - Validate Quorum Configuration
 - Validate Resource Status
 - Validate Service Principal Name
 - Validate Volume Consistency
- Inventory
 - Storage
 - List Fibre Channel Host Bus Adapters
 - List iSCSI Host Bus Adapters
 - List SAS Host Bus Adapters
 - System
 - List BIOS Information
 - List Environment Variables
 - List Memory Information
 - List Operating System Information
 - List Plug and Play Devices
 - List Running Processes
 - List Services Information
 - List Software Updates
 - List System Drivers
 - List System Information
 - List Unsigned Drivers
- Network
 - List Network Binding Order
 - Validate Cluster Network Configuration
 - Validate IP Configuration
 - Validate Network Communications
 - Validate Windows Firewall Configuration

- Storage
 - List Disks
 - List Potential Cluster Disks
 - **Validate CSV Network Bindings**
 - **Validate CSV Settings**
 - Validate Disk Access Latency
 - Validate Disk Arbitration
 - Validate Disk Failover
 - Validate File System
 - Validate Microsoft MPIO-Based Disks
 - Validate Multiple Arbitration
 - Validate SCSI device Vital Product Data (VPD)
 - Validate SCSI-3 Persistent Reservation
 - Validate Simultaneous Failover
 - **Validate Storage Spaces Persistent Reservation**
- System Configuration
 - Validate Active Directory Configuration
 - Validate All Drivers Signed
 - Validate Memory Dump Settings
 - **Validate Operating System Edition**
 - Validate Operating System Installation Option
 - Validate Operating System Version
 - Validate Required Services
 - Validate Same Processor Architecture
 - Validate Service Pack Levels
 - Validate Software Update Levels

Except for the Storage tests, all the tests can be run at any time because they aren't disruptive to the cluster.

Using the Validate a Configuration Wizard

Let's explore how to take advantage of the Validate a Configuration Wizard. Using the previous example of the problem related to adding

a node, let's say that the DNS server had the proper IP address and you can connect between the nodes outside the cluster. In this case, you can run the Validate a Configuration Wizard.

When you run the wizard, the Network/Validate Windows Firewall Configuration test fails. This test specifically looks at the Windows Firewall settings to ensure that port 3343, which is used by the cluster, hasn't been enabled. When this port is disabled, all communications on that port are blocked and you get errors in the Diagnostic channel.

Introducing the New Get-ClusterLog Command Switch

The Windows PowerShell command Get-ClusterLog lets you convert the events in a channel (e.g., FailoverClustering/Diagnostics) to a text file format. PowerShell will name the text file Cluster.log and place it in the C:\Windows\Cluster\Reports folder. If you run the command by itself, each node will have its own Cluster.log file. You can use switches to change this default behavior. Here are the switches, including the new -UseLocalTime switch:

- -Cluster <string>, where <string> is the name of the cluster you want to run the command against. This allows you to specify a remote cluster. If you omit the switch, it will run against the cluster you're currently on.
- -Node <string>, where <string> is the name of the node you want to run the command against. You use this command when you want to generate the Cluster.log file for a certain node only.
- -Destination <string>, where <string> is the folder to which you want to copy the Cluster.log files. If you include this switch, PowerShell will not only create a Cluster.log in each node's C:\Windows\Cluster\Reports folder but also copy all of the log files to the specified destination folder. This switch will add the node's name as part of the filename (e.g., Node1_Cluster.log, Node2_Cluster.log) for the log files copied to the destination folder. This way, each node's log files are easily identifiable.

- `-TimeSpan <uint32>`. You use this switch if you just want to get a log file that spans the last specified number of minutes, where `<uint32>` is that number (e.g., 5). This will give you a much smaller log file to review. You can use this switch if you're trying to reproduce an error. For example, you can reproduce the error you think might be occurring, then generate the log for the last 5 minutes to see if that's the case.
- `-UseLocalTime`. As mentioned previously, this switch is new in Server 2012. Clusters write all their information in GMT. For example, if you have a cluster that's in the GMT-5 time zone and your local time is 13:00 (1:00 P.M.), Cluster.log will show a time of 18:00 (6:00 P.M.) by default. With this switch, the local time is used, so the log will show a time of 13:00. When you use the `-UseLocalTime` switch, the times returned by the `Get-ClusterLog` command can easily be matched with the Event Log times.

Now that you know how to get Cluster.log files, it's time to learn how to read and search through them.

Reading Cluster.log Files

Reading Cluster.log files takes a long time to master, because they contain a lot of information that can be confusing. However, I'll give you some tips that can help you get started.

The first thing you need to understand is the anatomy of a Cluster.log file. Every entry has the same basic structure. Here's an entry for an IP address resource coming online:

```
000000bb8.0000001d4::2013/05/15-01:13:24.852
INFO [RES] IP Address <IP Address 1.1.1.1>:
Online: Opened object handle for netinterface
353c85ee-7ea7-4b2a-927d-1538dffcdcd
```

Let's break this entry down into smaller pieces to make better sense of it:

00000bb8. This is the process ID in hexadecimal notation. Typically, the process is the Resource Host System (RHS). You can see what resources the process is using by sorting or searching for the lines that include this process ID. This is useful when debugging an RHS dump if you have multiple files present. Each of these dumps is identified by a process ID, so knowing what the process ID is will ensure that you're working with the correct process dump. If you have a complete memory dump, there will be multiple RHS processes. Each is identified by the ID, so you can get to the correct one.

000001d4. This is the thread ID in hex notation. You can see what the thread is doing by sorting or searching for lines that include this thread ID. When you're debugging an RHS process that has 100 threads, you can jump right to the correct one using this ID.

2013/05/15-01:13:24.852. This is the date and time in GMT (unless the -UseLocalTime switch was used to generate the log). So if you're using GMT-5, the local time in this case is May 14, 2013, at 8:13 P.M. The time goes down to milliseconds.

INFO. This is the level of the entry. The level can be INFO (informational), WARN (warning), ERR (error), or DBG (debug). There are a few others, but these levels are what you'll see the majority of the time. Generally, a line with ERR in it indicates a problem with a resource. When you open a Cluster.log file after a failure, you can search for a specific level to try to get to the problem area quicker.

[RES] IP Address. This is the resource type. A resource will always identify its type in the log. With this information, you can more quickly follow the resource going online when there are multiple types of resources all coming online at the same time.

<IP Address 1.1.1.1>. This is the actual resource, as shown in Failover Cluster Manager.

Online: Opened object handle for netinterface 353c85ee-7ea7-4b2a-927d-1538dffcdcd. This is a description of what's going on with the resource. The resource is opening a handle to the network card driver in order to bind the IP address to it. If it fails here, it's most likely a

problem with the network card driver not accepting anything, which means it's bad. Alternatively, the network card might have died. Your next step would be to review the System event log entries to check for any network type events, such as the network going down or a card failing. With many of the descriptions, the more you see them, the more you'll understand what they mean. A description can be particularly helpful if it's describing the last action that occurred before a failure.

Searching Cluster.log Files

When reviewing Cluster.log files, it helps to search for keywords. Table 1 provides a list of keywords that I use when searching for resources. Note that you should type these keywords *exactly* as you see them. In other words, include the hyphen hyphen greater-than symbol (-->) and don't include any spaces.

Table 1: Keywords to Use When Searching for Resources

Keyword	Description
-->OnlinePending	This keyword appears in the log the second that Failover Cluster Manager displays Online Pending for a resource. This is where your search should start if you want to follow a resource coming online.
-->OfflinePending	This keyword appears in the log the second that Failover Cluster Manager displays Offline Pending for a resource. This is where your search should start if you want to follow a resource going offline.
-->Offline	This keyword appears in the log when Failover Cluster Manager displays Offline for a resource. So if you were following the resource, there's no need to look further. If this resource depends on another resource, that other resource could start its offline process first.
-->Online	This keyword appears in the log when Failover Cluster Manager displays Online for a resource. So if you were following the resource, there's no need to look further. If another resource depends on this resource, that other resource would not start its online process until this one completes.
-->ProcessingFailure	This keyword appears in the log when a resource just failed. If you find this line, you would want to start looking at previous entries to see what triggered the failure. Looking at entries after this event isn't necessary. Anytime a resource fails, you should still try to go through the normal offline process, even though you'll most likely get errors because the resource is unavailable.

You can also use these keywords to quickly determine how long a resource took to go offline or come online. For example, suppose that a group is taking longer than normal to come online. You can use --> OfflinePending as a starting point, then use --> Offline for all resources in the group. Alternatively, you can use --> OnlinePending followed by --> Online. For each resource, add up all the times to see how long it took to come online. After you've done that for all the resources, you can compare the resources' total times to see which resource took the longest amount of time. You can then review the entries in Cluster.log to determine why. For example, if a group took 30 seconds total to come online on one node and 3 minutes total to come online on another node, you should generate Cluster.log files for both nodes and compare them.

You can use the same keywords for groups, except that there must be a space after the greater-than symbol. For example, if a group goes offline, you would first use --> OfflinePending, followed by --> Offline. The only other difference between the resource entry and the group entry is that the group failure is --> Failed, whereas the resource failure is --> ProcessingFailure.

Putting It All Together

To see how all the information presented fits together, let's walk through solving a problem. Suppose that you have a two-node cluster configured with multiple file servers using different networks and a Fibre Channel connected SAN. Here's the setup for the networks:

- Cluster Network 1 = IP scheme 192.168.0.0/24
- Cluster Network 2 = IP scheme 1.0.0.0/8
- Cluster Network 3 = IP scheme 172.168.0.0/16

In the nodes' network connections, the network adapters are identified as:

- CORPNET = IP scheme 192.168.0.0/24
- MGMT = IP scheme 1.0.0.0/8
- BACKUP = IP scheme 172.168.0.0/16

FILESERVER1 is using Cluster Network 1, which is running on NODE1. FILESERVER2 is using Cluster Network 2, which is running on NODE2.

Let's say that there was a failure overnight and a file server group named FILESERVER2 was moved from NODE2 to NODE1. You need to find out why the failure occurred.

The first thing you do is open Failover Cluster Manager, right-click the FILESERVER2 group, and select Show Critical Events. This brings up the following events:

Event ID: 1069

Description: Cluster Resource 'IP Address 1.1.1.1' of type 'IP Address' in Clustered Role 'FILESERVER' failed.

Event ID: 1205

Description: The Cluster service failed to bring clustered service or application 'FILESERVER2' completely online or offline. One or more resources may be in a failed state.

The first event tells you that IP Address 1.1.1.1 experienced a failure. Therefore, you should right-click this resource in Failover Cluster Manager and choose Show Critical Events. You see the following events:

Event ID: 1077

Description: Health check for IP Interface 'IP Address 1.1.1.1' (address 1.1.1.1) failed (status is 1168). Run the Validate a Configuration wizard to ensure that the network adapter is functioning properly.

Event ID: 1069

Description: Cluster Resource 'IP Address 1.1.1.1' of type 'IP Address' in Clustered Role 'FILESERVER' failed.

Based on the description in first event (event 1077), you decide to use the Validate a Configuration Wizard. You want to run only the Network/Validate Network Communication test because that test will check the adapters and all network paths between the nodes.

After you run the Network/Validate Network Communication test, you check the test report. You don't see any errors or warnings, so you put it aside.

There are event channels you can review, so you go into the FailoverClustering/Operational channel, where you see this event:

Event ID: 1153

Description: The Cluster service is attempting to failover the clustered service or application 'FILESERVER2' from node 'NODE2' to node 'NODE1'

Because of this description, you go into the FailoverClustering/Diagnostics channel, where you see these events:

Event ID: 2051

Description: [RCM] rcm::RcmResource::HandleFailure:
(IP Address 1.1.1.1)

Event ID: 2051

Description: [RES] IP Address <IP Address 1.1.1.1>:
Failed to query properties of adapter id
F3EDD1C8-6984-82BC-498806B841CA, status 87.

Based on this information, you generate a Cluster.log file for this node. In the log, you search for --> ProcessingFailure and find these entries:

[RES] IP Address <IP Address 1.1.1.1>: IP Interface
3600A8C0 failed LooksAlive check, status 1168.
[RES] IP Address <IP Address 1.1.1.1>: IP Interface

```

3600A8C0 failed IsAlive check, status 1168.
[RHS] Resource IP Address 1.1.1.1 has indicated failure.
[RCM] Res IP Address 1.1.1.1: Online -> ProcessingFailure
( State Unknown )
[RCM] TransitionToState( IP Address 1.1.1.1)
Online-->ProcessingFailure.

```

A bit later in Cluster.log, you see the entries documenting that the group was being moved. This is a good indication that the entries found with the --> ProcessingFailure search are related to the problem that caused the group to be moved. Because of the errors seen in those entries, you know for sure that the IP address resource failed. To find out what the errors' status code means, you use the Net.exe command:

```
NET HELPMSG 1168
```

The command returns the message: *Element not found*. After looking more closely at the entries, it appears as though the actual problem might be with the network adapter. So, you run some hardware tests against the adapters and find that one adapter is faulty and not even showing up in Windows anymore. Replacing the faulty adapter is the course of action to fix the problem.

But there's still the question of why the Network/Validate Network Communication test results didn't show any errors when everything else did. This test checks all network adapters, going from one node to another, whether they're on the same network or not. It does this so that it knows all the routes it can take to get to the other nodes. So, there are some expected failures just because of the way the networks between the nodes are cabled or segmented.

You decide to look more closely at the test report. That's when you spot the output shown in Figure 2. You notice that NODE1 doesn't have a network adapter defined as MGMT. This is basically saying the same thing as the events, which is that NODE1 has two networks and

Result	Source Network Interface	Destination Network Interface	Same Cluster Network
Success	NODE1 - CORPNET	NODE2 - CORPNET	Yes
Failed	NODE1 - CORPNET	NODE2 - MGMT	No
Failed	NODE1 - CORPNET	NODE2 - BACKUP	No
Failed	NODE1 - BACKUP	NODE2 - CORPNET	No
Failed	NODE1 - BACKUP	NODE2 - MGMT	No
Success	NODE1 - BACKUP	NODE2 - BACKUP	Yes

Result	Source Network Interface	Destination Network Interface	Same Cluster Network
Success	NODE2 - CORPNET	NODE1 - CORPNET	Yes
Failed	NODE2 - CORPNET	NODE1 - BACKUP	No
Failed	NODE2 - MGMT	NODE1 - CORPNET	No
Failed	NODE2 - MGMT	NODE1 - BACKUP	No
Failed	NODE2 - BACKUP	NODE1 - CORPNET	No
Success	NODE2 - BACKUP	NODE1 - BACKUP	Yes

Figure 2
Network/
Validate Network
Communication Test
Results

NODE2 has three networks. So, the lesson here is that you need to do more than just look at the errors or warnings at the top of the report. You also need to look at the test results.

Get to the Root of the Problem

Troubleshooting a cluster is like troubleshooting just about anything. There are different ways to troubleshoot and multiple things to look at in order to get to a problem’s root cause. I presented one way to get to the root cause, and I hope you’re able to use it when troubleshooting problems in your clusters. For more information pertaining to failover clustering, check out the [Ask the Core Team](#) blog site and the [Clustering and High Availability](#) blog site. ■

Enabling the Next-Generation File Server with SMB 3.0

New features expand your file-level storage options

Every organization uses Server Message Block (SMB) in some form to access storage. It might be to access logon scripts, to access and use software-installation media, or for users to access their documents and MP3 collections. But what SMB hasn't been used for is a file-level protocol (in which the client doesn't directly access the disk blocks but instead is served files) for enterprise applications to access remote storage. When it comes to communicating with storage for enterprise workloads, block-level technologies (in which the server can communicate directly with disk blocks) such as iSCSI and Fibre Channel (and maybe NFS for non-Windows workloads) have been top on the list.

Consider the difference between a user accessing a document on a file share and an enterprise application storing its database on a file share. For a user editing a Microsoft PowerPoint document from an SMB share, portions of the document are cached locally, and occasionally the user clicks Save. If the SMB file-server experiences a problem such as rebooting, or if it's clustered and the file share moves to another node in the cluster, the user loses the handle and lock to the file—but without any real impact. The next time the user clicks



John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Microsoft Virtualization Secrets* (Wiley).



Save, everything is re-established and no harm is done. Now consider Hyper-V storing a virtual machine (VM) on an SMB file share that experiences a problem. The file share moves to another node in the cluster. First, the Hyper-V server waits for the TCP timeout before realizing that the original connection has gone. This could mean 30 seconds of pause to the VM. But Hyper-V has also now lost its handles and locks on the virtual hard disk (VHD), which is a major problem. Whereas user documents might be used for a few hours, enterprise services such as a VM or database expects handles on files to be available for months without interruption.

Fortunately, SMB 3.0 addresses this issue, and many more. For [Windows Server 2012](#), Microsoft wanted to make SMB a file-level storage protocol that could be used for crucial enterprise workloads such as Microsoft Hyper-V and SQL Server. To make this shift, some major changes to the SMB protocol were required.

Enabling Transparent Failover

If SMB is being used to house enterprise data such as VMs and SQL Server databases, then it's unlikely to be used on a standalone file server. Rather, it will be part of a cluster, to provide high availability. For a clustered file service, a single cluster node typically mounts the LUN that contains the shared file system and offers the share to SMB clients. If that node fails, then another node in the cluster mounts the LUN and offers the file share. However, the SMB client then loses its handles and locks.

SMB Transparent Failover provides protection from a node failure. It does so by enabling a share to move between nodes in a manner that is completely transparent to the SMB clients, maintaining any locks and handles that exist as well as maintaining the state of the SMB connection.

The state of the SMB connection is maintained over three entities: the SMB client, the SMB server, and the disk that holds the data. SMB Transparent Failover ensures that enough context exists to bring the

SMB connection state back to an alternate node if a node fails, allowing SMB activities to continue without the risk of error.

However, even with SMB Transparent Failover, there can still be a pause to I/O. The LUN must be mounted on a new node in the cluster. But the Failover Clustering team has done a huge amount of work around optimizing the dismount and mount of a LUN to ensure that it never takes more than 25 seconds. That sounds like a long time, but it's the absolute worst-case scenario, involving large numbers of LUNs and tens of thousands of handles. For most common scenarios, the time would be only a couple seconds. And enterprise services such as Hyper-V and SQL Server can handle an I/O operation of 25 seconds without error.

Another possible cause of interruption to I/O is the SMB client noticing that the SMB server is unavailable. In a typical planned scenario (e.g., a node rebooting because it's being patched), the server notifies clients, which can then take the appropriate actions. But if a node crashes, there is no client notification. Rather, the client sits and waits for TCP timeout before taking action to re-establish connectivity—a waste of resources. Although an SMB client might have no idea that the node it's talking to in the cluster has crashed, the other nodes in the cluster know within a second, thanks to the various IsAlive messages that are sent between nodes.

This knowledge is leveraged by the new Witness Service, available in Windows Server 2012. The Witness Service essentially allows another node in the cluster to act as a witness for the SMB client. If the node that the client is talking to fails, the witness node notifies the SMB client, allowing the client to connect to another node and minimizing the service interruption to a couple seconds. The conversation looks something like the following (but in 1s and 0s and with less personality):

SMB Client to Server A: “I want to establish a connection, Server A.”

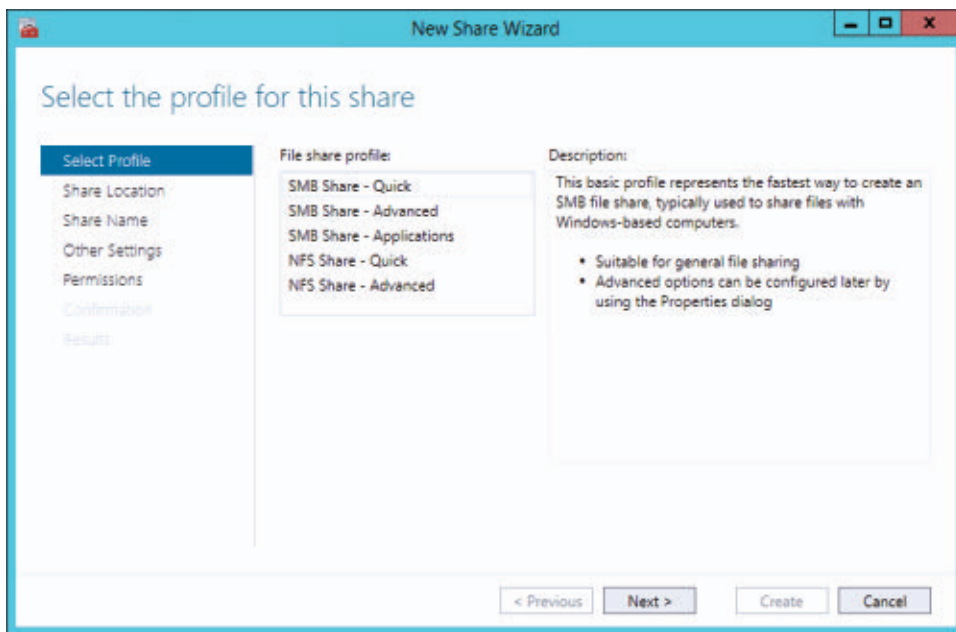
Server A: “The connection is established. Also, I am part of a cluster. Servers B, C, and D are also in the cluster.”

SMB Client to Server B: “Server B, I have established an SMB connection to Server A. Can you watch Server A and notify me if it fails?”

Server B: “Yes. Have a nice day.”

The good news is that you don’t need to do anything special to take advantage of SMB Transparent Failover or the Witness Service. When you create a new share on a Windows Server 2012 failover cluster, SMB Transparent Failover is enabled automatically. A wizard guides the process of creating a new share in a Windows Server 2012 file server cluster. The first decision is which type of share you are creating. The answer simply helps to set some default options for the file share, as shown in Figure 1. But for all SMB Share types, the *Enable continuous availability* setting is enabled, as shown in Figure 2.

Figure 1
Creating Supported
Share Types



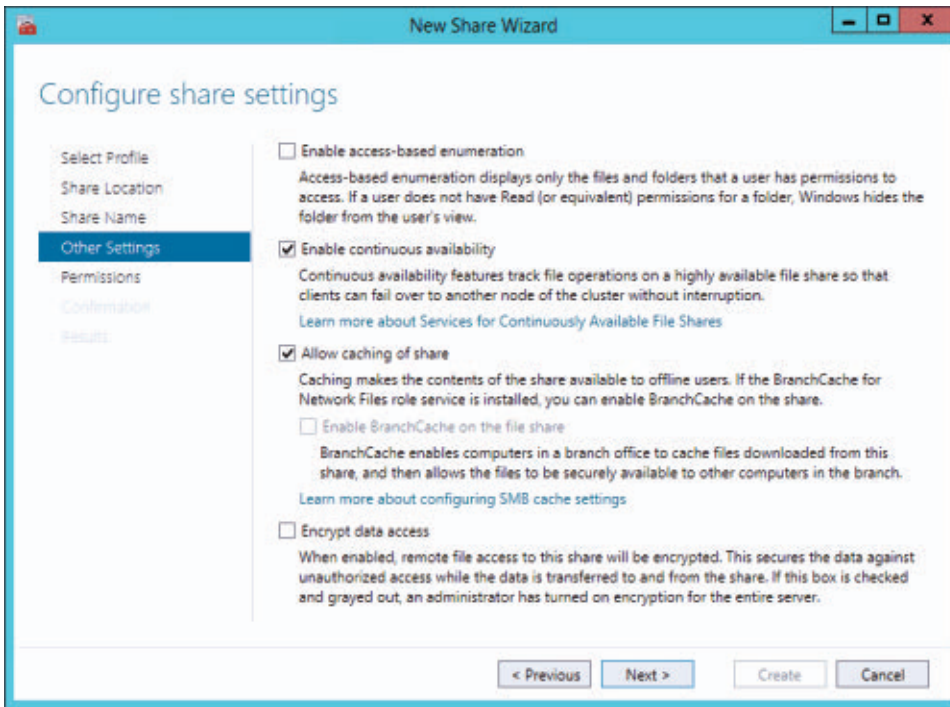


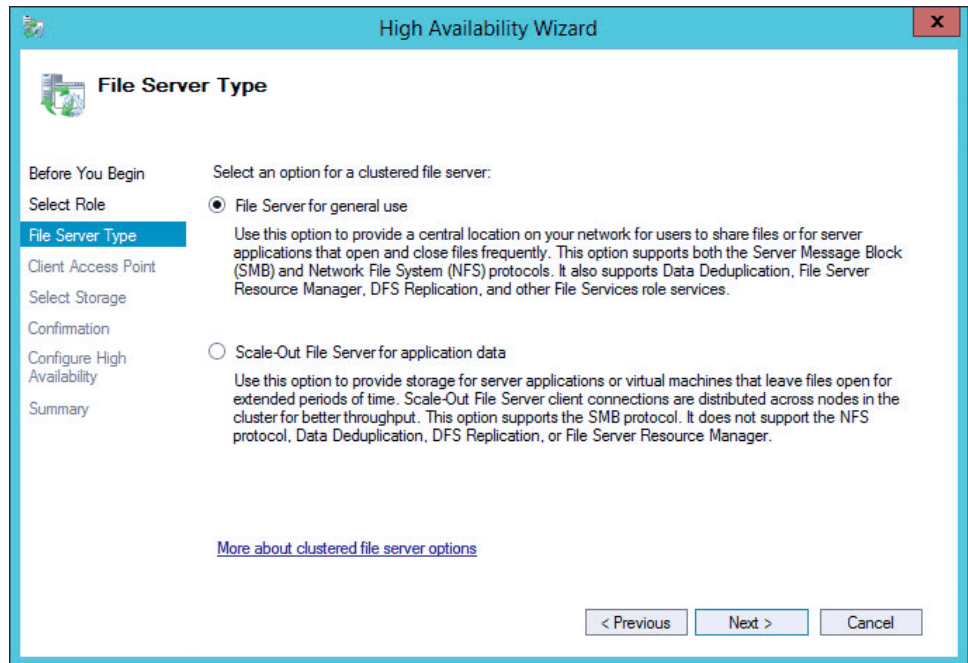
Figure 2
Available Options
for a Share

SMB Active/Active Configuration

I discussed the necessity of a brief I/O pause as the shared LUN is moved between nodes. You might be familiar with this as a challenge for Windows Server 2008 Hyper-V when moving VMs between nodes. The problem stems from the fact that NTFS is a shared-nothing file system and can't be accessed concurrently by multiple OS instances without the risk of corruption. This problem was solved with the introduction of cluster shared volume (CSV) support in Windows Server 2008 R2. CSV allows all nodes in a cluster to read and write to a set of LUNs simultaneously, using some clever techniques and removing the need to dismount and mount LUNs between nodes.

Windows Server 2012 extends the use of CSVs to a specific type of file server, namely the new Scale-Out File Server. This option is targeted for use only when sharing application data such as SQL Server

Figure 3
Creating a Scale-Out
File Server on a CSV



databases and Hyper-V VMs. The traditional style of a general-use file server is still available for non-application data, as shown in Figure 3.

When you choose the option to create a Scale-Out File Server, you must also choose a CSV to use as storage when shares are subsequently created within the file server. Because this storage is available to all nodes in the cluster, all those nodes also host the file share. Therefore, SMB client connections are distributed over all the nodes instead of just one. If a node fails, no work is involved in moving the LUNs, offering an even better experience and reducing interruption in operations to almost zero. This reduction is crucial for the application-server workloads at which this Scale-Out File Server is targeted.

The use of Scale-Out File Servers offers an additional benefit. Typically, when a general-use File Server is created, you must give the new cluster file server a NetBIOS name and unique IP address as part of the configuration. That IP address must be hosted by whichever cluster

node is currently hosting the file server. With Scale-Out File Servers, all nodes in the cluster offer the file service. Therefore, no additional IP addresses are required. Instead, the IP addresses of the nodes in the cluster are used via the configured Distributed Network Name (DNN).

I should point out that although all nodes in the cluster offer the same file service—and therefore shares—with the Scale-Out File Server, any single SMB client will connect to only one node at any one time. Essentially, when the SMB client initiates connections, it initially receives a list of all the IP addresses for the hosts in the cluster. The client picks one with which to initiate the SMB session and then uses only that node, unless the node experiences a problem. If that happens, the client converses with an alternate node, except when leveraging the Witness Service.

Protecting Against Connection Failure: SMB Multichannel

SMB Transparent Failover and SMB active/active configuration are great technologies that help protect against interruptions caused by a node failure. But there are other types of failure, such as a connection failure. To counteract this type of issue, you can use technologies such as Microsoft Multipath I/O (MPIO), which provides multiple paths from server to storage. SMB 3.0 introduces SMB Multichannel, which allows an SMB client to establish multiple connections for a single session, providing protection from a single connection failure and boosting performance.

Like most SMB 3.0 features, SMB Multichannel happens automatically. After the initial SMB connection is established, the SMB client looks for additional paths to the SMB server. If multiple network connections are present, those additional paths are used. The use of SMB Multichannel is apparent when monitoring a file copy operation, because only one connection's worth of bandwidth is used initially but doubles as the second connection is established, continues to increase with the third connection, and so on. If a connection fails, other connections continue the SMB channel without interruption.

To determine whether SMB Multichannel is in effect on a server, use the Get-SMBConnection Windows PowerShell cmdlet, which shows the SMB connections to an SMB share. In the output that Figure 4 shows, I can see that I have only one connection to my server.

```
PS C:\> Get-SmbConnection
```

ServerName	ShareName	UserName	Credential	Dialect	NumOpens
savdalhv01	software	SAVILLTECH\john	SAVILLTECH.NET\john	3.00	1

Figure 4

Listing All the Current
SMB Connections

This output indicates that there is only one usable path between the SMB client and the SMB server. If I run the Get-SmbMultichannelConnection cmdlet from the client, the output shows all the possible paths over which the server can accept connections, as shown in Figure 5.

```
PS C:\> Get-SmbMultichannelConnection
```

Server Name	Selected	Client IP	Server IP
savdalhv01	True	192.168.1.15	192.168.1.30
savdalhv01	True	192.168.1.15	10.1.3.1
savdalhv01	True	192.168.1.15	10.1.2.1

Figure 5

Identifying Possible
Paths for the SMB
Multichannel

However, this list is generated by a “lazy” check and does not mean that a path can actually be created between the client and server IP addresses 10.1.3.1 and 10.1.2.1.

To confirm which path is actually being used between the client and the server, I can look at the TCP connections to remote port 445, which is used for SMB. This confirms that I am using the one path that can be used: remote address 192.168.1.30, as Figure 6 shows.

```
PS C:\> Get-NetTCPConnection -remoteport 445
```

LocalAddress	LocalPort	RemoteAddress	RemotePort	State	AppliedSetting
192.168.1.15	55071	192.168.1.30	445	Established	Internet

Figure 6

Finding Actual
Connections Used for
SMB

A common question, if your SMB client connects to an SMB share that’s hosted on an active/active cluster, is whether those multiple connections occur to different nodes in the cluster. The answer is no. The SMB client receives a single IP address for one node in the

cluster, and all connections are to that node. All SMB sessions for that cluster from one SMB client will always go to the same node in the cluster. Remember, this isn't a problem because a highly available cluster typically has hundreds if not thousands of connecting SMB clients. The load will be distributed fairly evenly throughout the cluster.

Maximizing Bandwidth: Receive Side Scaling and Remote Direct Memory Access

The final aspect of SMB 3.0 that I want to focus on relates to the larger network-connection pipes in today's data center. Many data centers have shifted from 1Gbps to 10Gbps. But as data centers adopt 10Gbps, the processor in the server becomes a performance bottleneck. A single TCP connection can be processed by only one processor core, which can't handle 10Gbps and typically restricts the bandwidth. This is where Receive Side Scaling (RSS) comes into play. With RSS, a single network interface is split into multiple receiving connections, each of which can be serviced by a separate processing core. Therefore, the full bandwidth can be utilized. Most modern server network adapters automatically support RSS. To determine whether your hardware supports RSS, run the `Get-SmbMultichannelConnection` cmdlet, as shown in Figure 7.

Note that this output shows the number "4" for both `CurrentChannels` and `MaxChannels`. This is the default for Windows Server 2012 when leveraging RSS-capable network cards. If you then look at the SMB connections from the server, which Figure 8 shows, you'll see that four separate connections are established for the IP address that SMB uses, confirming that RSS is in action.

```
PS C:\> Get-SmbMultichannelConnection | fl
ServerName       : savdalhv01
Selected         : True
Failed           : False
FailureCount     : 0
ClientInterfaceIndex : 14
ClientRSSCapable : True
ClientRdmaCapable : False
ClientLinkSpeed  : 1 Gbps
ClientIpAddress  : 192.168.1.31
ServerInterfaceIndex : 14
ServerRSSCapable : True
ServerRdmaCapable : False
ServerLinkSpeed  : 1 Gbps
ServerIpAddress  : 192.168.1.30
MaxChannels      : 4
CurrentChannels  : 4
```

Figure 7
Viewing the SMB
Multichannel
Configuration

```
PS C:\> get-nettcpconnection -remoteport 445
```

LocalAddress	LocalPort	RemoteAddress	RemotePort	State	AppliedSetting
-----	-----	-----	-----	-----	-----
:::1	49156	:::1	445	Established	Datacenter
192.168.1.30	58617	192.168.1.35	445	Established	Datacenter
192.168.1.30	58615	192.168.1.35	445	Established	Datacenter
192.168.1.30	58614	192.168.1.35	445	Established	Datacenter
192.168.1.30	56715	192.168.1.35	445	Established	Datacenter

Figure 8

Identifying
Current SMB Client
Connections on a
Server

You might wonder why an RSS-capable network interface is split into four connections by default. (You can confirm this default by using the `Get-SmbClientConfiguration` PowerShell cmdlet to look at the SMB configuration. The first line of the output shows the connection count per RSS network interface.) You can change this value, but the number wasn't picked randomly. Microsoft went through much testing on 10Gbps connections and found that four connections produces the most gain; more than four connections brings diminishing returns. However, if you have connections larger than 10Gbps, then increasing this value might benefit you.

Remote Direct Memory Access (RDMA) is another technology that brings high throughput performance and minimizes server load. Network adapters that support RDMA can bypass most of the network stack to communicate directly, avoiding load on the host servers. The `Get-SmbMultichannelConnection` cmdlet that I referred to earlier will show whether the network adapter supports RDMA. During the initial SMB connection initialization, a check is performed to determine whether both ends of the connection support RDMA. If they do, the connection switches to RDMA. Again, no manual setup is required.

A Powerful Solution

SMB 3.0 is used only between OSs that support SMB 3.0, namely Windows Server 2012 and [Windows 8](#). For other OSs, a negotiation is performed and the highest common version of SMB supported is used. For example, if a Windows 7 machine connects to a Windows Server 2012 file server, then SMB 2.1 is used because that's the highest version that Windows 7 supports.

The primary driver for most of the changes in SMB 3.0 was the desire to make SMB an enterprise-application protocol. That is certainly where you'll see the biggest benefit to SMB. But there are still benefits for regular clients, such as Windows 8 clients. (SMB 3.0 is unavailable for OSs earlier than Windows 8 and Windows Server 2012.) For example, the new SMB encryption capability removes the need for complicated public key infrastructures (PKIs) to achieve protection. SMB 3.0, along with many other Windows Server 2012 storage changes, puts the new OS on the map as a powerful storage solution and gives customers even more choice. ■



FREE Newsletters!

Not your average Newsletters!
subscribe today at windowsitpro.com/manage-newsletters

WinInfo Daily UPDATE
 Paul Thurrott covers the entire Windows universe with reviews, commentary, analysis, and tips. Delivered daily.

Windows IT Pro UPDATE
 Windows industry news, products, FAQs, tips, and resources for IT professionals. Delivered weekly.

Cloud & Virtualization UPDATE
 Get the latest news, blogs and analysis to help you determine your organization's cloud and virtualization strategy. Delivered weekly.

Exchange and Outlook UPDATE
 News, strategies, products, and developments in Exchange Server and Outlook messaging. Delivered weekly.

Security UPDATE
 Learn about Windows security risks, attacks, and how to fix or avoid them. Includes security alerts! Delivered bi-weekly.

Dev Pro UPDATE
 Topics for Microsoft platform developers: ASP.NET, .NET Framework, Silverlight, mobile, and SQL Server development. Delivered weekly.

SQL Server Pro UPDATE
 The latest news, products, and developments for SQL Server DBAs and developers. Delivered weekly.

SharePoint Pro UPDATE
 SharePoint for IT professionals and developers – weekly tips, news, and how-to's. Delivered weekly.

Windows IT Pro

Windows Server 2012: Making DHCP Highly Available

New DHCP failover feature simplifies the process



Orin Thomas

is a contributing editor for *Windows IT Pro* and a Windows Security MVP. He has authored or coauthored more than a dozen books for Microsoft Press.

Email



Blog



It's no secret that DHCP is a critical component in network infrastructure. On most networks, client computers use DHCP to receive their IP address information. The problem is, unlike DNS, making DHCP highly available hasn't always been a straightforward task. If a DHCP server fails and isn't returned to service quickly, clients will be unable to access the network because they won't have valid IP addresses. Unless you have a monitoring solution in place, it's likely that the first time you'll know about a DHCP failure is when an increasing number of users with IP addresses in the Automatic Private IP Addressing (APIPA) range call the service desk.

In previous OSs, such as Windows Server 2008 R2 and Windows Server 2003, you have two options to make DHCP highly available:

- You can put the DHCP server on a failover cluster, with the configuration information stored on shared storage.
- You can configure split scopes. This involves carving up a normal scope so that 80 percent of the addresses in the lease are hosted on the DHCP server most likely to respond to client traffic on a particular subnet. The remaining addresses in the lease are hosted on a DHCP server on a remote subnet and are used by clients only when the DHCP server with 80 percent of the addresses isn't available.

[Windows Server 2012](#) simplifies and improves DHCP availability by introducing a DHCP failover feature to the DHCP role service. DHCP failover lets you provide a highly available DHCP service without having to configure split scopes or deploy a failover cluster. After I give you more details about this new feature, I'll show you how to configure it.

Understanding DHCP Failover

DHCP failover involves configuring two Server 2012 computers with the DHCP role service installed as a pair. This pair can provide a highly available DNS using one of the following techniques:

- *Load balance mode.* The load balance mode (which is sometimes referred to as the load sharing mode in the Microsoft documentation) is the default method of configuring DHCP failover. When you configure two DHCP servers in load balance mode, each server will serve IP addresses from the same scope in such a way that duplicate addresses aren't issued. Address leases from the scope are issued by each server in a load balanced manner. If one DHCP server fails, the other DHCP server will continue to lease addresses until the first DHCP server returns to service. Figure 1 shows a DHCP scope configured to use the load balance mode.

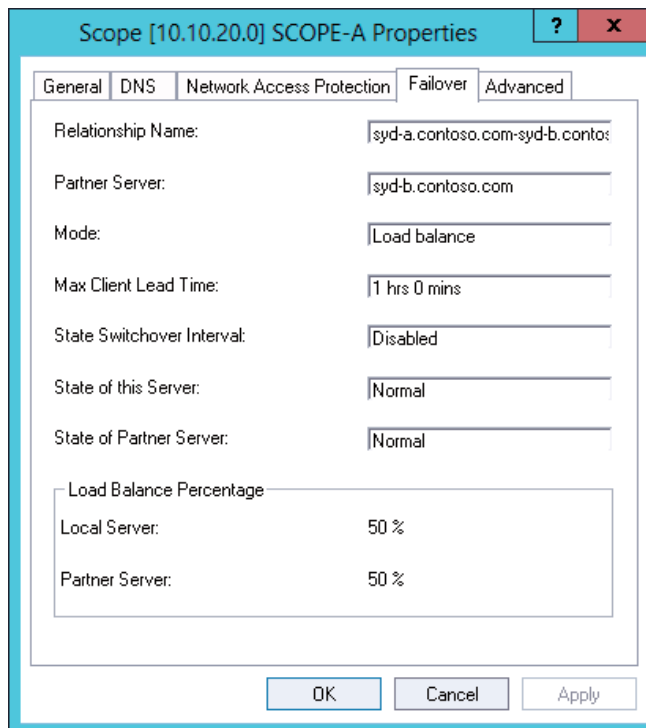
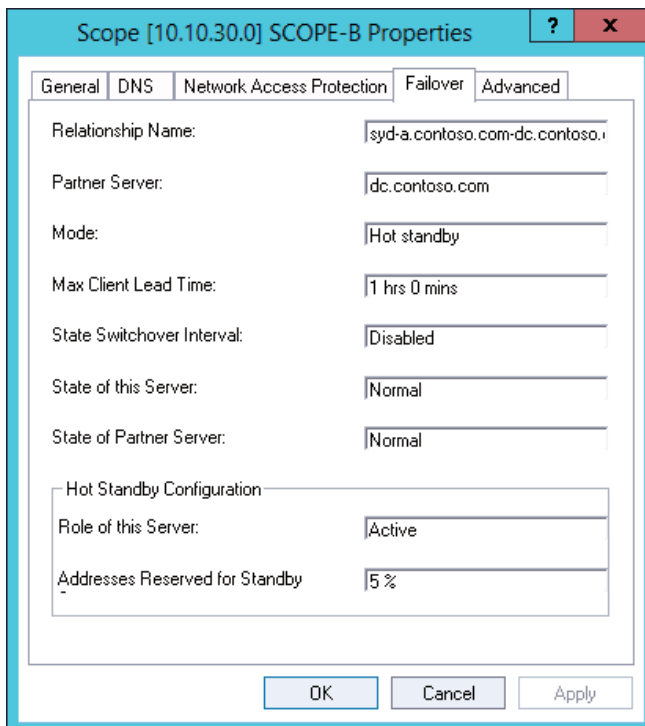


Figure 1
Reviewing the Properties of a DHCP Scope Configured to Use the Load Balance Mode

- *Hot standby mode.* When you configure two servers with the DHCP role installed in hot standby mode, the servers operate in a failover relationship. The active server leases IP addresses and configuration information to clients. The secondary server only performs this function in the event that the primary isn't available. Figure 2 shows a DHCP scope configured to use the hot standby mode.

Figure 2
Reviewing the
Properties of a DHCP
Scope Configured to
Use the Hot Standby
Mode



Configuring DHCP Failover

DHCP failover involves setting up a partnership between two DHCP servers. Only two DHCP servers can participate in a partnership, but you can configure multiple partnerships between DHCP servers. For example, you can configure DHCP-ONE and DHCP-TWO as partners, DHCP-TWO and DHCP-THREE as partners, and DHCP-ONE and DHCP-THREE as partners. An individual DHCP scope, however, can only be used with one partnership. For

example, you can configure SCOPE-ALPHA as highly available on servers DHCP-ONE and DHCP-TWO, but this scope can't also be present on DHCP-THREE.

To configure DHCP failover, perform the following steps:

1. Install the DHCP role on two separate servers running Server 2012 that are members of the same Active Directory (AD) domain.
2. Ensure that the DHCP role on each server is authorized in AD.
3. Create the relevant scopes on the first DHCP server.
4. Click the scope for which you want to configure failover. On the Action menu, click Configure Failover.
5. On the *Introduction to DHCP Failover* page of the Configure Failover wizard, verify that the scope you selected is present and click Next.
6. On the *Specify the partner server to use for failover* page, click Add Server. As Figure 3 shows, the Add Server dialog box will list all the Server 2012 computers running the DHCP role service that have been authorized in the domain. Select the DHCP server you want to use as the partner and click OK.
7. On the *Specify the partner server to use for failover* page, click Next.

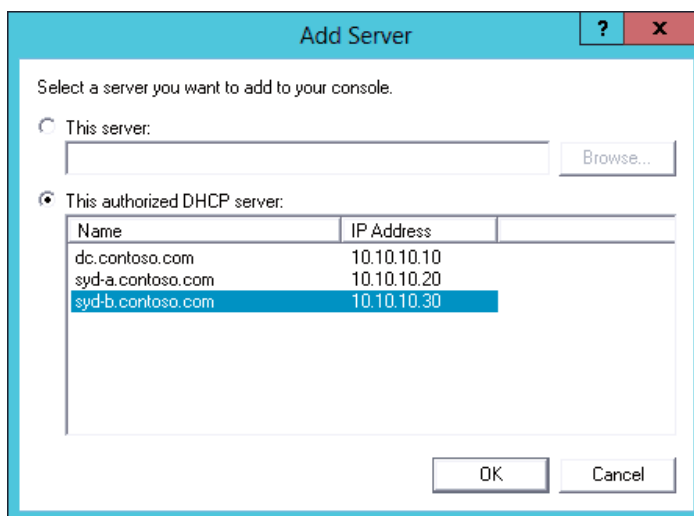


Figure 3
Selecting the DHCP
Server to Use as a
Partner

8. On the *Create a new failover relationship* page, select either *Load balance* or *Hot standby* in the Mode drop-down list.
9. If you're configuring the server to use the load balance mode, specify the weight assigned to each server. The default is that each server shares an equal load, as shown in Figure 4.

Figure 4
Configuring the DHCP
Pair to Use the Load
Balance Mode

Configure Failover

Create a new failover relationship

Create a new failover relationship with partner syd-b.contoso.com

Relationship Name: syd-a.contoso.com-syd-b.contoso.com

Maximum Client Lead Time: 1 hours 0 minutes

Mode: Load balance

Load Balance Percentage

Local Server: 50%

Partner Server: 50%

☐ State Switchover Interval: 60 minutes

☒ Enable Message Authentication

Shared Secret:

< Back Next > Cancel

If you're configuring the server to use the hot standby mode, specify the role of the partner server (which can be set to Active or Standby) and the percentage of addresses in the scope reserved for the standby server, as shown in Figure 5.

10. If desired, configure the State Switchover Interval option. The setting determines the length of time before the standby begins leasing addresses to clients on the network.

Configure Failover

Create a new failover relationship

Create a new failover relationship with partner syd-b.contoso.com

Relationship Name:

Maximum Client Lead Time: hours minutes

Mode:

Hot Standby Configuration

Role of Partner Server:

Addresses reserved for standby server: %

☐ State Switchover Interval: minutes

☒ Enable Message Authentication

Shared Secret:

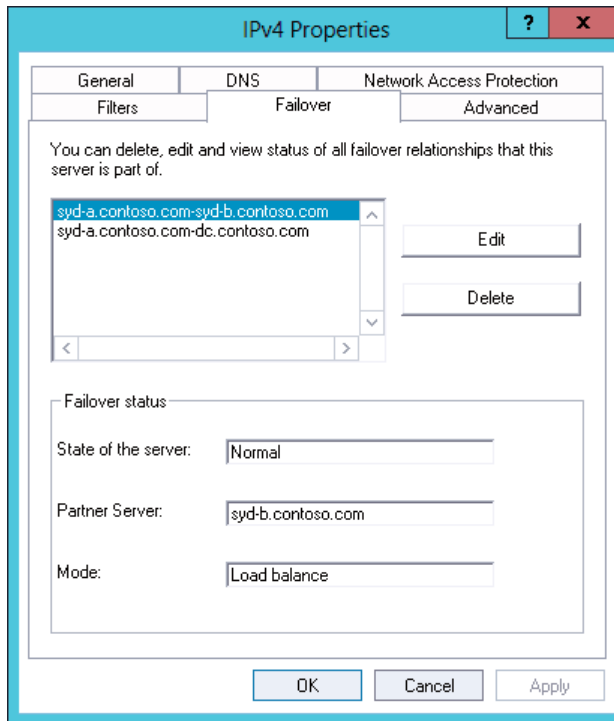
Figure 5

Configuring the DHCP Pair to Use the Hot Standby Mode

11. Choose a shared secret. This allows you to pair the DHCP servers. Click Next.
12. On the final page, click Finish.

You can configure only one type of failover relationship between two DHCP servers. So, if you configure a load balance relationship between DHCP-ONE and DHCP-TWO, all scopes configured for DHCP failover will need to use the load balance mode. If you configure a relationship between DHCP-ONE and DHCP-THREE, that relationship can use a different failover method. You can view the relationships that a DHCP server has on the Failover tab of either the IPv6 Properties or IPv4 Properties dialog box, as Figure 6 shows.

Figure 6
 Reviewing the DHCP
 Server's DHCP Failover
 Relationships



Create a Highly Available DHCP Solution with Minimal Work

DHCP failover in Server 2012 provides a highly available DHCP solution without requiring you to configure split scopes or a failover cluster. In most situations, using the default load balanced DHCP failover configuration will be suitable. You can configure multiple relationships between different DHCP servers, but you can only make a scope highly available on a single relationship. ■



JOHAN ARWIDMARK



SEAN DEUBY



CHANDER DHALL



MARY JO FOLEY



DAN HOLME



MARK MINASI



MICHAEL OTEY



MARK RUSSINOVICH



ROD TRENT



PAUL THURROTT



For more than 13 years, IT/Dev Connections has been the premier training event for developers and IT professionals. IT/Dev Connections provides in-depth training on the technology platforms you're currently using, real-world solutions that will give you the competitive edge, and expert insight into how to plan for and implement the latest technologies. With more than 175 sessions to choose from, the conference offers training on HTML5, ASP.NET, Exchange, SQL Server, Windows OS, Windows Server, SharePoint, Visual Studio, Office 365, business intelligence, cloud, and all types of development.

BROUGHT TO YOU BY ▼



Go to **devconnections.com** to register

▼ Schedule at a Glance *(subject to change)*

MONDAY

7:30am	Registration Opens
9:00am – 4:00pm	Pre-conference workshops

TUESDAY

7:00am – 5:00pm	Registration Open
7:30am – 8:30am	Breakfast
8:30am – 9:30am	Keynote
9:30am – 10:00am	Break
10:00am – 11:15am	Sessions
11:15am – 11:45am	Break
11:45am – 1:00pm	Sessions
1:00pm – 2:30pm	Lunch
2:30pm – 3:45pm	Sessions
3:45pm – 4:15pm	Break
4:15pm – 5:30pm	Sessions
5:30pm – 7:30pm	Welcome Reception

WEDNESDAY

7:00am – 5:00pm	Registration Open
7:30am – 8:30am	Breakfast
8:30am – 9:30am	Keynote
9:30am – 10:00am	Break
10:00am – 11:15am	Sessions
11:15am – 11:45am	Break
11:45am – 1:00pm	Sessions
1:00pm – 2:30pm	Lunch
2:30pm – 3:45pm	Sessions
3:45pm – 4:15pm	Break
4:15pm – 5:30pm	Sessions

THURSDAY

7:00am – 5:00pm	Registration Open
7:30am – 8:30am	Breakfast
8:30am – 9:30am	Keynote
9:30am – 10:00am	Break
10:00am – 11:15am	Sessions
11:15am – 11:45am	Break
11:45am – 1:00pm	Sessions
1:00pm – 2:30pm	Lunch
2:30pm – 3:45pm	Sessions
3:45pm – 4:15pm	Break
4:15pm – 5:30pm	Sessions

FRIDAY

7:30am	Registration Opens
9:00am – 4:00pm	Post-conference workshops

Make Connections the conference you bring your whole team to this year!

NEW LOOK!
same great conference you know and love!

For more than 13 years, Connections has been the go-to Microsoft training conference for IT Professionals and Developers.

IT/Dev Connections is the only in-person training conference you need to attend this year. Join thousands of Microsoft developers, DBAs and IT Professionals and stay on top of the latest in:

- ASP.NET
- HTML5
- Visual Studio
- SQL Server
- Windows Server
- SharePoint
- Exchange
- and much, much more.

Register with the
**All Access
VIP PASS
and Save!**



CONFERENCE AT A GLANCE

3 full days of educational sessions, with pre- and post-conference workshops

Train with 100+ Microsoft & industry experts

More than 175 in-depth technical, how-to sessions

Exciting and entertaining networking opportunities

Discover new products and services in the partner expo hall

Ask YOUR pressing technical strategy questions, hear diverse opinions from our expert speakers and develop your company's plan for the future!

HIGHLY ANTICIPATED SPEAKERS



MARY JO FOLEY
Microsoft Expert
AllAboutMicrosoft.com



MIGUEL DE ICAZA
CTO, Co-founder of the
GNOME and MONO projects
Xamarin



MARK MINASI
Senior Contributing Editor
Windows IT Pro



MARK RUSSINOVICH
Senior Contributing Editor
Windows IT Pro



PAUL THURROTT
Senior Technical Analyst
Windows IT Pro

MORE EXPERT SPEAKERS

This year's conference speakers were chosen out of more than 500 stellar session entries, guaranteeing the best-of-the-best are speaking at Connections 2013.



JOHAN ARWIDMARK
Knowledge Factory



ITZIK BEN-GAN
SQL Server Pro



JES SCHULTZ BORLAND
Brent Ozar Unlimited



ANDREW CONNELL
Critical Path Training



BRIAN DESMOND
Moran Technologies



SEAN DEUBY
Windows IT Pro



CHANDER DHALL
Microsoft MVP



TIM FORD
Spectrum Health



SCOT HILLIER
SharePoint MVP



DAN HOLME
Microsoft MVP



STEVE JONES
SQL Server Central



MICHAEL OTEY
SQL Server Pro



BRENT OZAR
Brent Ozar Unlimited



JEREMIAH PESCHKA
Microsoft MVP



TONY REDMOND
Windows IT Pro



PAUL ROBICHAUX
Windows IT Pro



LORYAN STRANT
Office 365 MVP



JEREMY THAKE
Microsoft MVP



ORIN THOMAS
Windows IT Pro



ROD TRENT
Windows IT Pro



RANDY WILLIAMS
AvePoint



NATHAN WINTERS
Microsoft

...and many more!

See website for updated speakers and sessions.

(Speakers are subject to change.)

175+ sessions

*Pick and choose
from more than 175+
sessions to optimize
your training
experience.*

HTML5 | Visual Studio | ASP.NET
Windows | Exchange | Cloud
SQL Server | SharePoint

Sessions subject to change.

*Go to www.devconnections.com to see the most
up-to-date schedule.*

SQL Server

- How Active Directory Affects SQL Server
- Manage Your Shop with CMS and Policy Based Management
- Inside the Query Optimizer
- Data Internals Deep Dive
- T-SQL Querying and Query Tuning Enhancements in the Latest Major Releases of SQL Server
- Efficient Interval Management in SQL Server
- Practical Uses of Window Functions
- Use Dynamic Management Views to Diagnose SQL Server Performance Issues
- Hardware 201: Selecting and Sizing Database Hardware for OLTP Performance
- SQL Server Transaction Log Internals
- SQL Server Index Internals
- SQL Server 2012 in a Highly Available World
- Table Indexing for the .NET Developer
- SQL Server Table Partitioning from the Ground Up
- Using Power View and Hadoop to Unlock Hidden Markets
- Build Your Own SQL Server Private Cloud
- Surviving Your Peak Database Load
- Troubleshooting SQL Server with SysInternals Tools

- Code-Less Securing of SQL Server
- Shortcuts to Productivity in SQL Server Management Studio
- From Zero to Hero: A Case Study in Reducing Extremely High I/O on a SQL Server System
- SQL Server Optimization: Tuning the Hardware Subsystems
- What DBAs Need to Know About Hekaton
- Using FullText Search with Office Documents and PDFs
- Encryption in SQL Server
- Collecting and Analyzing File and Wait Statistics
- Maximizing Plan Reuse
- Practical Performance Monitoring in SQL Server
- Windows Azure SQL Database Troubleshooting and Query Tuning
- Windows Azure SQL Database for the DBA
- Make Your Queries Fly With Columnstore Indexes
- Using BIML as an SSIS Design Patterns Engine
- Hacking the SSIS 2012 Catalog
- SSIS Design Patterns
- Reduce, Reuse, Recycle: Automating Your BI Framework
- BI Security Best Practices
- From Reporting Services Rookie to Rock Star
- Maximizing SSIS Package Performance
- Cleaning Up Dirty Data in SSIS
- Choosing Between SSAS 2012 Multidimensional and Tabular
- PowerPivot to SSAS 2012 Tabular
- Managing SQL Server Performance with Extended Events
- Leveraging the Plan Cache for Performance Tuning
- Improve the Performance of Your T-SQL Queries by Changing Your Habits
- Creating SSRS Reports Efficiently Through Best Practices
- Manage SQL Server Efficiently with PowerShell Remoting
- Manage SQL Server 2012 on Windows Server Core with PowerShell
- TempDB Performance Troubleshooting and Optimizing
- Page Latches for Mere Mortals

- Creating Highly Performant, Scalable Websites: From SPA to Backend
- Architecting Device-driven Mobile Web Solutions
- Building Secured, Scalable, Low-Latency Web Applications with the Windows Azure Platform
- Writing Next Generation JavaScript with TypeScript
- Creating Data-Driven Mobile Web Apps with ASP.NET MVC and jQuery Mobile
- Touch-Enabled and Data Connected Sites in Hours, Not Weeks
- Unit Testing ASP.NET MVC
- Software Gardening
- Branches and Merges Are Bears, Oh My!
- Exploring Domain-Driven Design Implementation Patterns in .NET
- Behavior-Driven Development: Turning User Stories into Executable Specifications
- CQRS: Crack for Architecture Addicts?
- Get More Bang for Your Windows Azure Buck!
- Grokking Caching
- From Developer to Architect: 10 Things You Must Know
- Everything You Need to Know About Trends in Application Development
- Linq to Everything
- Asynchronous Programming with Async and Await
- Build your first Angular Web Application
- Building Games for Windows 8 – Using GameMaker
- Best Practices for Building Windows Phone and Windows 8 Applications
- Domain-Driven Design, CQRS, and Event-Sourcing for the Busy Developer
- Do's and Don'ts of Software Projects
- Fast Facts of Social Network Programming
- Creating Data-Driven HTML5 Applications
- Building End-to-End Web Apps Using TypeScript
- jQuery Fundamentals
- Automating Windows Azure from the Command Line
- Caching in Azure: There's More to That Than Azure Caching
- Debugging the Web with Fiddler
- IIS for Developers
- Doing It Right: Continuous Delivery Doesn't Have to Suck

- Building a Windows 8 App from Scratch
- Git for Visual Studio Developers
- From Manual Testing to Automation with Visual Studio ALM
- New Features in Visual Studio 2013 and TFS 2013
- Strategies for Refactoring and Testing Legacy Code
- Connecting the Dots: Using HTML5, jQuery, and Web API Together
- Advanced Debugging with WinDbg and SOS
- Task and Data Parallelism: Real-World Examples
- .NET Garbage Collection Performance Tips
- How to (Remote) Control Office 365 with Windows Azure
- Using Async in Your Mobile Apps
- A .NET Developer's Guide to Mobile Apps
- Cloud Data for the Everyday Developer
- Building Solutions in the Cloud with Apps for Office
- Developing Professional Solutions for Office 2013 and Outlook
- Doing Modern Web in the Enterprise
- Web Performance Optimization for Modern Web Applications
- Debugging and Testing JavaScript in Today's Browsers
- Developing Neural Networks with Visual Studio
- Introduction to iPhone Programming with C#, .NET, and Xamarin.iOS
- Introduction to Android Programming with C#, .NET, and Xamarin.Android
- Introduction to Mobile Web with HTML5
- Essential Typescript
- Unit Testing Web Development
- Customizing the SharePoint 2013 user interface with JavaScript
- JavaScript for Windows 8 Developers, Part 1 and 2
- JavaScript Testing - An Introduction
- Building Cross-Platform Mobile Applications with PhoneGap, Part 1 and Part 2
- Developing with the SharePoint 2013 App Model
- Creating Line-of-Business Apps in HTML5 and MVC/Web API
- Everyday Bootstrap
- Simplify Your API: Creating Maintainable and Discoverable Code

- Introduction to PowerShell for the Anxious IT Pro
- SharePoint Performance: Best Practices from the Field
- Who Says You Can't Do Records Management in SharePoint?
- Top 10 New ECM Features in SharePoint 2013
- Dan Holme's SharePoint 2013 MasterClass: SharePoint Installation and Configuration, From Bare Metal to Farm
- Implementing End-to-End SharePoint Governance
- Best Practices for Role-Based Management of Users, Groups, Permissions, Service Accounts, and Administrative Delegation
- Developing Search Applications in SharePoint 2013
- Developers Approach to Social Applications with SharePoint 2013
- The Only Way to Go is Up! Upgrading to SharePoint 2013
- Optimizing and Accelerating Your SharePoint Farm
- 0 to 60: Apps for Office and SharePoint
- Migrating SharePoint Solutions to Apps for SharePoint
- What Options Do Non-Developers Have in SharePoint 2013?
- Create Powerful SharePoint Designer 2013 Workflows in Office 365 and On-Premises
- Office 365: Introduction to SharePoint Online Development
- Surfacing Your Azure External Data Using BCS in SharePoint 2013 with Alerts
- Data Visualization with SharePoint and SQL Server
- Extending the Business Process Management Features of Office 365

- Exchange Online: Real-World Migration Challenges and Solutions
- Notes from the Field: Running a 500,000-Mailbox On-Premises Exchange Server Deployment
- Exchange Server 2013 Site Resiliency
- Managed Availability: Ensuring the End User Experience
- Data Loss Prevention in the Real World
- Hybrid and SSO Deployment with the New Office 365 (Wave 15)
- The Tao Of Exchange Server 2013 Sizing
- Virtualizing Exchange Server 2013: Why Not?
- From Zero to Hero: PowerShell for Exchange Server Boot Camp
- Exchange Server 2013 Unified Messaging Deep Dive
- Better Together: Integrating Exchange Server 2013 and Lync Server 2013
- Migrate to Modern Public Folders the Worry-Free Way
- Troubleshooting Modern Public Folders: A DIY Guide
- Apples to Apples: Comparing Office 365 to the Competition
- How Does Microsoft Secure My Email with Office 365?
- CAS 2013 – Why It Is 3 Better Than CAS 2010 and 6 Better Than 2007
- Building a Hybrid Configuration with Exchange Server 2013 in (Less Than) 75 Minutes
- How-to: Load Balancing Exchange Server 2013
- Exchange ActiveSync: Taming the Beast
- Exchange Server 2013 Backup, Restore, and Recovery
- High Availability in Exchange: A Recipe for Success?

- Managing Third-Party Updates with System Center 2012 Configuration Manager SP1
- Migrating from Configuration Manager 2007 to Configuration Manager 2012
- Developing Hydration Kits – IT Pro Automation at Its Best!
- A Geek's Guide to USMT 5.0
- Configuration Manager 2012 SP1 OS Deployment
- Using Windows Azure Infrastructure as a Service as Your Data Center
- What's New in Windows Server 2012 Hyper-V
- The WHY of Configuration Manager
- Hierarchy Simplification with Configuration Manager 2012
- Deploying and Managing Virtual Applications and Settings with System Center Configuration Manager and MDOP
- Deploying and Managing Virtual Applications and Settings with Active Directory Domain Services and MDOP
- Smoothing the Kinks for a Seamless User Experience with Microsoft UE-V
- Deploying Your Office in the Cloud with Office 365
- Windows Server 2012 Advanced Troubleshooting Workshop
- Troubleshooting Group Policy in Windows Server 2012
- Managing Public Cloud Infrastructure with PowerShell
- Manage Server 2012 Like a Pro or, Better, Like an Evil Overlord!
- Windows "Next:" Will Blue Make You Blue?
- AppLocker: Your Solution for Application Smackdown!
- Deploy Office 2010 or Office 2013 Using Group Policy (It CAN Be Done!)
- Windows Installer Survival Guide for System Center Configuration Manager Admins
- System Center Configuration Manager Software Update Zen
- Using Orchestrator to Integrate with Azure IaaS
- Become an Orchestrator Master
- Hyper-V Best Practices
- Microsoft Windows PowerShell Remoting In-Depth
- State-Based Administration of the Modern Enterprise
- Configuration Manager for UNIX and Mac — Myths and Realities
- Windows Intune Overview
- Managing Devices in the Cloud with Windows Intune

Mandalay Bay Resort & Casino

Network with your colleagues at Mandalay Bay Resort & Casino! There's so much to do, you'll never have to leave this 4-star resort!

HOTEL ACCOMMODATIONS

Mandalay Bay Resort & Casino
3950 Las Vegas Blvd. South, Las Vegas, NV
SPACE IS LIMITED so reserve your room early.
Call: 877-632-9001 and reference IT/Dev Connections
Room Block Rates Expire September 15, 2013

ATTIRE

The recommended dress for the conference is casual and comfortable. Please bring along a sweater or jacket, as the ballrooms can get cool with the hotel's air conditioning.

TAX DEDUCTIONS

Your attendance to a DevConnections conference may be tax deductible. Visit www.irs.ustreas.gov. Look for topic 513 — Educational Expenses. You may be able to deduct the conference fee if you undertake to (1) maintain or improve your skills required in your present job; (2) fulfill an employment condition mandated by your employer to keep your salary, status, or job.

GROUP DISCOUNTS

Register individuals from one company at the same time and receive a group discount (10% off registration. Not to be combined with other discounts or offers).



Registration & Cancellation Policy: Registrations are not confirmed until payment is received. Cancellations before August 2, 2013, must be received in writing and will be refunded minus a \$100 processing fee. After August 2, 2013, cancellations and no shows are liable for full registration; it can be transferred to the next Conference within 12 months or to another person. You may transfer this registration to a colleague by notifying us before the start of the event. Please inform us if you have any special needs or dietary restrictions when you register. The Conference Producers reserve the right to cancel the conference by refunding the registration fee. Producers can substitute speakers and topics and cancel sessions without notice or obligation. Updates will be posted on our website at www.DevConnections.com.

Notes & Policies: Tape recording, video recording and photography are not allowed at any session. Conference producers will be taking candid pictures of events and reserve the right to reproduce. By attending this conference you agree to this policy. Microsoft, Microsoft .NET, ASP.NET, Visual Studio, Microsoft SQL Server, Exchange, SharePoint and Windows are either trademarks or registered trademarks of Microsoft Corporation. All other trademarks are property of their owners.

Conference Registration

Full Conference Registration Includes Keynote on October 1, 2013, through Closing Session October 3, 2013.

Name	Discount Code
Company	Title
Street Address (Required to ship materials)	
City/State/Postal Code	
Country	
Phone	Fax
E-Mail (required)	

Online:
devconnections.com

Email:
Info@devconnections.com

Phone:
888.899.0130

Fax:
800.766.5367

Mail:
Penton Media
DevConnections
24654 Network Place
Chicago, IL 60673

Check the conference track you are registering for.
NOTE: you can attend any of the co-located conference tracks for no additional charge.

- ☐ Dev Connections
- ☐ Windows Connections
- ☐ SQL Server Connections
- ☐ SharePoint Connections
- ☐ Exchange Connections
- ☐ ALL ACCESS VIP PASS..... \$2,695
- ☐ BASIC REGISTRATION \$1,695
- ☐ Pre-Conference Workshops
Monday, September 30, 2013..... \$499
- ☐ Post-Conference Workshops
Friday, October 4, 2013..... \$499

TOTAL

Payment Information:

☐ CHECK
*(payable to Penton Media) All payments must be in US currency.
Checks must be drawn on a US bank.*

☐ CREDIT CARD:
☐ VISA ☐ MASTERCARD ☐ AMEX

--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

Credit Card No.

--	--	--	--	--	--

Expiration Date

Cardholder's Signature

Cardholder's Name (print)

PowerShell Basics: Variables

Variables are an integral part of how PowerShell works

Variables are a fundamental part of [Windows PowerShell](#). They're quite different from the variables in Cmd.exe. In Cmd.exe, all variables are environment variables, which you primarily access with the Set command. The environment variables can store only strings of text. (You can store a number in an environment variable, but it's up to the program that's reading the environment variable to interpret it as a number.)

The variables in PowerShell aren't environment variables but rather native PowerShell variables. They can store much more than just text. In fact, PowerShell variables store *objects* (specifically, Microsoft .NET Framework objects). For example, a PowerShell variable can contain a String object or a number object, such as an Int (integer) object. Objects provide an extraordinary amount of flexibility.

Creating Variables

In PowerShell, variable names start with the \$ character. You can assign a value to a variable using the assignment operator, which is the = character. You can create a variable by simply assigning it a value. For example, the command

```
$myName = "Ferb"
```

creates a variable named *\$myName* and assigns it a string value. The double quotes (" ") indicate that a string value is being assigned to the variable.



Bill Stewart

is a scripting guru who works for Indian Health Service in Albuquerque, New Mexico. He's a contributing editor for *Windows IT Pro* and a moderator for Microsoft's Scripting Guys forum. He offers free tools on his website.



Email



Website

As I mentioned previously, PowerShell variables are really objects. In simple terms, objects can contain data (properties) and operations you can perform on the data (methods). In this example, the `$myName` variable is really a String object. As with other objects, the String object has both properties and methods. For example, the `Length` property of a String object tells you the number of characters in the string, and the `ToUpper` method gives you a copy of the string converted to uppercase. You can access both properties and methods using a dot (.) after the variable name. Properties don't use parentheses (), but methods do. For example, the command

```
$myName.Length
```

returns a value of 4 because the variable's value (Ferb) is four characters long. The command

```
$myName.ToUpper()
```

returns *FERB*.

Discovering an Object's Type, Properties, and Methods

The properties and methods that an object can use depend on the object's type. For example, a String object has different properties and methods than an Int object. You can get a variable's object type by calling its `GetType` method like this:

```
$myName.GetType()
```

Figure 1

Determining a Variable's Object Type by Calling Its `GetType` Method

As Figure 1 shows, the `$myName` variable contains a String object.

PS C:\> \$myName.GetType()			
IsPublic	IsSerial	Name	BaseType
-----	-----	-----	-----
True	True	String	System.Object

(In this introductory discussion, I won't talk about the `IsPublic`, `IsSerial`, and `BaseType` columns.)

Besides using the `GetType` method to find out the kind of object a variable contains, you can also use the `Get-Member` cmdlet to see what properties and methods are available. For example, if you run the command

```
Get-Member -InputObject $myName
```

you'll find that 35 properties and methods are available. Figure 2 shows a few of them.

```
PS C:\> $myName | Get-Member
```

Name	MemberType	Definition
Clone	Method	System.Object Clone()
CompareTo	Method	int CompareTo(System.Object value), ...
Contains	Method	bool Contains(string value)
.		
.		
Length	Property	System.Int32 Length {get;}

Figure 2

Discovering the Available Properties and Methods with the `Get-Member` Cmdlet

Introducing Collections

As you've seen, variables can store a single object (e.g., a `String` or `Int` object). Variables can also store multiple objects, which are referred to as a *collection* or an *array*. For example, the command

```
$items = Get-ChildItem
```

uses the `Get-ChildItem` cmdlet to retrieve the collection of file system objects in the current directory (i.e., the directory from which you're running the command) and stores that collection in the variable `$items`.

Introducing Variable Interpolation

When you include a variable's name inside a double-quoted string, PowerShell replaces the variable's name with its value in the string. This is called *variable interpolation*. For example, if you run the commands

```
$myName = "Ferb"  
"Hello, $myName"
```

you'll receive the result *Hello, Ferb*.

If the variable you're expanding isn't a string, PowerShell will do its best to coerce the variable's value into a string representation. In addition, PowerShell doesn't perform variable interpolation for single-quoted strings, so you can use single-quoted strings when you don't want PowerShell to replace variables in a string.

A common problem when using variable interpolation is when you want to include an object's property (or the result of an object's method) in the string. Using the standard dot notation to retrieve the property doesn't quite work as expected. For example, the following command

```
"$myName is $myName.Length characters"
```

returns the incorrect result of *Ferb is Ferb.Length characters*.

To work around this problem, PowerShell provides the *subexpression operator*, `$()`, which you can use within the string to get the desired result. For example, the command

```
"$myName is $($myName.Length) characters"
```

returns the correct result of *Ferb is 4 characters*.

In general, if PowerShell isn't replacing a variable in a double-quoted string like you expect it to, you can put the variable inside `$()` to work around the problem.

Exploring Automatic and Preference Variables

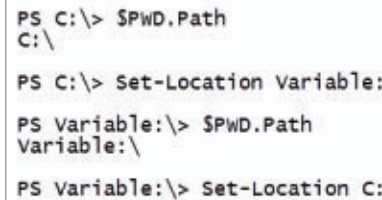
PowerShell provides the Variable: drive, which lets you access the variables in your PowerShell session. The Variable: drive is just like a file system drive (e.g., C), except you're accessing variables instead

of file system items. For example, this command lists all variables in your current session:

Get-ChildItem Variable:

When you run this command, you'll see a list of variables and their values. Initially, these variables are all *automatic variables* (which store something about PowerShell's state) and *preference variables* (which store PowerShell user preferences). You can't change the values of automatic variables, but you can change the values of preference variables.

Automatic variables tell you something about PowerShell's current state. For example, the \$PWD automatic variable contains PowerShell's current location. Consider the commands and their output in Figure 3.



```
PS C:\> $PWD.Path
C:\

PS C:\> Set-Location Variable:

PS Variable:\> $PWD.Path
Variable:\

PS Variable:\> Set-Location C:
```

Figure 3

Using the \$PWD
Automatic Variable

The \$PWD variable contains an object, and the first command retrieves the value of its Path property. At this point, the value is C:\. The next command uses the Set-Location cmdlet to change the current location to the Variable: drive. Note that the Path property of the \$PWD variable reflects the change automatically (hence the term automatic). But just for verification, the third command again retrieves the value of the \$PWD variable's Path property, which is Variable:\ this time. The last command sets the current location back to the C drive.

Preference variables let you change a user preference within PowerShell. One of the most common preference variables is \$ErrorActionPreference, which lets you configure how PowerShell should respond to non-terminating errors. (Non-terminating errors don't prevent the cmdlet from continuing.) By default, \$ErrorActionPreference is set to Continue, which means PowerShell will output the non-terminating error and the cmdlet will continue running. Sometimes you might

want to have the cmdlet stop as soon as it encounters an error, in which case you'd run the command:

```
$ErrorActionPreference = "Stop"
```

Other times, you might not care about non-terminating errors, so you just want to suppress them altogether with this command:

```
$ErrorActionPreference = "SilentlyContinue"
```

Note that using the `$ErrorActionPreference` variable has the same effect as using the `-ErrorAction` cmdlet parameter. The difference is that the `-ErrorAction` cmdlet parameter affects only a single cmdlet, whereas the `$ErrorActionPreference` variable affects all cmdlets.

Exploring Environment Variables

PowerShell provides the `Env:` drive, which lets you access environment variables. In PowerShell, environment variables get copied from the parent process (i.e., the program that started the current PowerShell session). Typically, the initial values of the environment variables are the same as those in Control Panel. (You can use a PowerShell profile script to change the initial values of environment variables so that they don't match the values in Control Panel, but that's beyond the scope of this discussion.)

To view all environment variables in the current PowerShell session, you can run the command:

```
Get-ChildItem Env:
```

This is equivalent to running the `Set` command in `Cmd.exe`.

In `Cmd.exe`, you can surround a string with `%` characters (e.g., `%ALLUSERSPROFILE%`), which tells `Cmd.exe` to replace the variable's name with its value. PowerShell doesn't use the `%` characters

to get the values of environment variables. In PowerShell, you can access an environment variable's value two ways. First, you can access it from the Env: drive directly using the syntax `$Env:name`, where *name* is the environment variable's name. For example, if you want to find out the value of `ALLUSERSPROFILE`, you'd run the command

```
$Env:ALLUSERSPROFILE
```

which returns *C:\ProgramData*. Alternatively, you can use the `Get-Item` cmdlet to retrieve the value of an environment variable from the Env: drive. In this case, you don't use the `$` character, as shown here:

```
Get-Item Env:ALLUSERSPROFILE
```

The first syntax (`$Env:name`) is the most common and works when using variable interpolation in double-quoted strings. For example, the command

```
"The ALLUSERSPROFILE variable is $Env:ALLUSERSPROFILE"
```

returns *The ALLUSERSPROFILE variable is C:\ProgramData*.

Getting Help with Variables

The PowerShell Help system provides quite a bit of information about variables. I recommend that you take a look at the following Help topics: `about_Variables`, `about_Automatic_Variables`, `about_Preference_Variables`, and `about_Environment_Variables`. To read these Help topics online, you can follow this syntax: `help topic`, where *topic* is the topic you want to display. For example, if you want to read the `about_Variables` topic, you'd run the command:

```
help about_Variables
```

If you're running PowerShell 3.0 and you get an error message when trying to display a Help topic, you'll need to download the Help topics first. To do this, start PowerShell as an administrator by right-clicking the PowerShell shortcut icon and choosing *Run as administrator*. At the PowerShell prompt, type the command:

```
Update-Help
```

Your computer must have a working Internet connection to be able to download the Help topics.

Get a Grip on PowerShell Variables

Understanding PowerShell variables is important because variables are such an integral part of how PowerShell works. Although this brief introduction doesn't explain all there is to know about PowerShell variables, the information presented here provides the essentials you need to understand. ■

IPv6 Support in Windows 8 and Windows Server 2012

Avoid IPv6 problems

Microsoft has a strong history of supporting IPv6, with solutions dating as far back as the days of Windows 2000. Microsoft is continuing its support for IPv6 in [Windows 8](#) and [Windows Server 2012](#). I'll provide you with an overview of the IPv6 capabilities in these latest Microsoft OSs and highlight some potential areas of concern. I won't be going into too much detail about what IPv6 is or how it works. For additional information about IPv6, see the [Learning Path](#).

IPv6 Overview

You've probably heard experts' warnings that the world is running out of new IPv4 address blocks. The American Registry for Internet Numbers (ARIN) has about 44 million remaining free IP addresses it can hand out in blocks. Although that might seem like a lot, it's only 0.01 percent of the theoretical maximum number of IPv4 addresses. The situation is worse in the Middle East, Europe, Central Asia, and Asia Pacific. For example, there are fewer than 16 million remaining free IPv4 addresses that can be handed out by the Reseaux IP Europeens Network Coordination Centre (RIPE NCC) and Asia-Pacific Network Information Centre (APNIC), which are the Regional Internet Registries (RIRs) for those regions.

If a company is in the market for more IPv4 addresses, it can still get them from either its ISP or an RIR directly if it qualifies. That situation will likely remain for some time. However, IPv6 is the future of the Internet. Many new products and services support only IPv6 or use IPv6 by default. For example, IPv6 is the default in Windows 7



John Howie

is the Chief Operating Officer of the Cloud Security Alliance. He is also a visiting professor at the University of Arizona and at Edinburgh Napier University.



Email



Twitter



LinkedIn

and later and Windows Server 2008 and later. In addition, some organizations (including the U.S. government) are mandating that all new computing products and services being obtained support IPv6. If you plan to use a new product that defaults to or exclusively uses IPv6 or if you want to do business with an organization that's mandating its use, your networks will need to support IPv6. As a result, you need to start planning for IPv6 if you haven't already done so.

The key differences between IPv6 and IPv4 are twofold. First, IPv6 addresses are 128 bits in length, which is four times longer than IPv4 addresses. Second, the addressing scheme used in IPv6 is very different from IPv4. In IPv4, you have several classes of addresses, special addresses for nonpublic use, and some other edge cases that were added as new Internet products and technologies were developed. IPv6 cleaned a lot of that up, so the addressing is easier to understand. For specific details, check out the [Learning Path](#).

During the development of IPv6, it became clear that systems would need to support both IPv4 and IPv6 concurrently as well as provide a means for IPv6-only systems to access IPv4-only systems. I'll discuss how Windows 8 and Server 2012 meet these needs next. During the development of IPv6, it also became clear that there was a need to provide a means to transition from IPv4 to IPv6 without replacing all the existing network hardware. This is where some concerns exist for security professionals, which I'll discuss later in the "[Security Concerns](#)" section.

Windows 8 and Server 2012 IPv6 Support Out of the Box

Windows 8 and Server 2012 support IPv6 out of the box. Server 2012 further supports IPv6 by providing:

- Support for the Dynamic Host Configuration Protocol for IPv6 (DHCPv6).
- IPv6 addresses in the DNS server.
- Transition technologies such as Network Address Translation for IPv6 to IPv4 (NAT64) and DNS for IPv6 to IPv4 (DNS64). These

two technologies are used in Server 2012's DirectAccess feature, which heavily uses IPv6.

You can't remove IPv6 support from Windows 8 and Server 2012, but you can disable it. In fact, I highly recommend disabling IPv6 in your organization until you're ready to configure and use it. You can disable it in corporate environments by editing the registry, using Group Policy with policy scripts you've created, or using Microsoft *Fix it* scripts that must be run on each machine on which you want to disable IPv6. You can also simply unbind IPv6 from the physical adapters, but IPv6 will still be running and can still be used to connect to IPv6 sites over IPv4. You can find more details in the Microsoft Support article "[How to disable IP version 6 or its specific components in Windows.](#)"

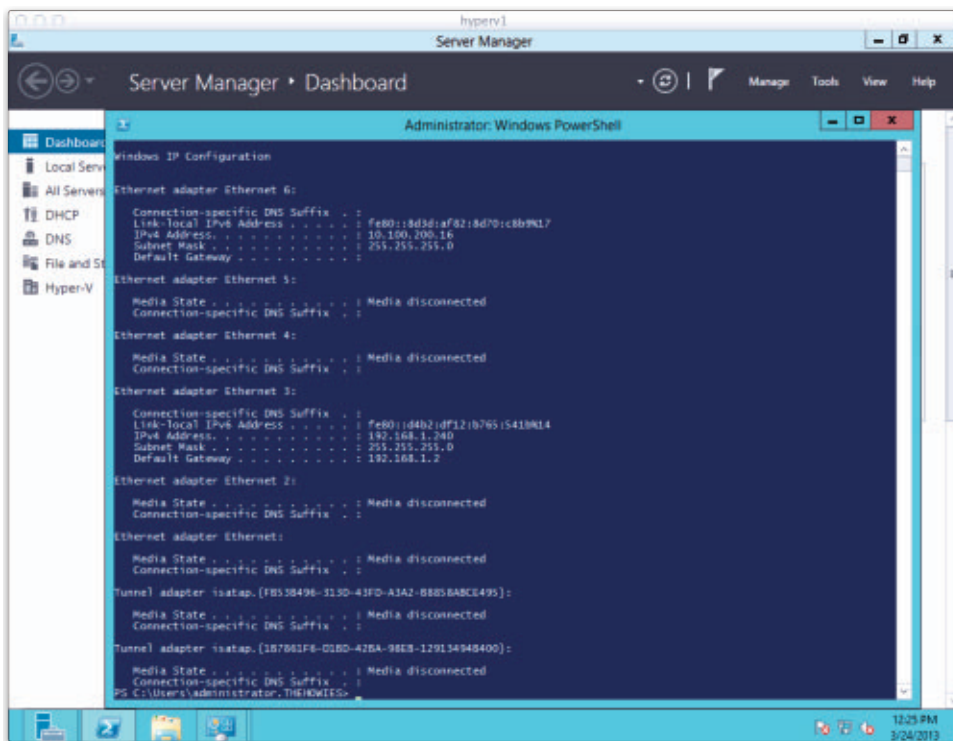
Unlike previous Windows OS versions, Windows 8 and Server 2012 don't give you the option to specify the network configuration when installing fresh copies of them. When the OSs are installed, Windows will auto-configure IPv4 and IPv6 addresses using a variety of technologies. You'll likely be familiar with some of these technologies but not others. A word of caution is that Windows 7 and later and Server 2008 and later will do their best to obtain an IPv6 address, even if you ask them not to. I'll explain this further in the next section.

There are also some areas of potential concern in that not all IPv6 support in Windows is standards compliant. Although this noncompliance probably won't cause you any problems, you need to be aware of it. In many cases, you can use the Netsh utility or Windows PowerShell to force the Windows OS to be standards compliant.

IPv6 Address Configuration

Windows 8 and Server 2012 use a variety of techniques to obtain IPv6 addresses for each adapter present on the machine. Even when Windows is unable to obtain routable IPv6 addresses, it configures interfaces with link-local IPv6 addresses, as shown in Figure 1.

Figure 1
Default Network
Configuration with
Link-Local IPv6
Addresses



There is no practical way to stop the allocation of link-local addresses—nor should you want to disable them, because link-local addresses are used for communications between hosts and between hosts and routers. By default, Windows won't use link-local IPv6 addresses to communicate, but it's important to understand that they can be used and that they can be used by default if you really want Windows to use them (or if you make a significant number of mistakes in how you configure Windows networking).

If an IPv6-ready networking infrastructure isn't configured, Windows 8 and Server 2012 will still be able to use IPv6 and configure IPv6 addresses in certain situations:

- **Situation 1: Home users with public IP addresses.** In this situation, Windows will try to establish a connection using the IPv6 transition technology named Teredo. Teredo will work only if the

Windows machine isn't domain-joined and has UDP access to the Internet, with no firewall-blocking packets.

- Situation 2: Home users with public IP addresses when Teredo fails. In this situation, Windows will use another IPv6 transition technology named 6to4. It requires only a publicly routable IP address.
- Situation 3: Windows can resolve the name using the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) by means of DNS or name broadcasts. In this situation, Windows will assume that the host is an ISATAP server that's capable of accepting IPv6 packets encapsulated in IPv4 packets, delivering them to IPv6 hosts, encapsulating replies, and sending the replies back. ISATAP works in both domain-joined and non-domain-joined environments. It also works in RFC 1918 nonroutable IP address environments.

If you want domain-joined Windows 8 and Server 2012 systems to use IPv6, you'll likely want to assign predetermined IPv6 addresses to each system, especially if they're Server 2012 systems. At a minimum, you need to provide the IPv6 address allocated to each system. Optionally, you can also provide the IPv6 address of the default gateway that each system should use and the DNS server's IPv6 addresses. Providing this information is optional because Windows can get it from other sources. In the case of the default gateway, Windows can participate in router solicitation and listen for router advertisements to learn the IPv6 addresses of routers. It can also use IPv4 addresses to communicate with DNS servers. Using DNS over IPv4 assumes that your DNS servers are used to store IPv6 addresses of hosts in your organization and are capable of making recursive queries to other servers to get addresses for hosts outside your organization.

If you have numerous Windows 8 systems or a Server 2012 system that you really don't care which IPv6 address it gets, you might be tempted to use DHCPv6. However, I highly recommend resisting that temptation. To understand why, you need to understand how the IPv6 address configuration process works.

When Windows 8 or Server 2012 starts up, it sends out router solicitation requests to find IPv6-capable routers. Routers will respond to router solicitation requests and will periodically send out router advertisements with their address. Routers can also provide additional information, such as the addresses of DNS servers and domain search suffixes. When Windows receives a response or hears an advertisement for an adapter whose IPv6 address hasn't been configured, it will use the information provided by the router to configure an IPv6 address—even if the router asks it not to.

IPv6 routers use two flags to tell an IPv6 client what to do with the information they provide. The first flag is the Managed Address Configuration flag (or simply the *m* flag). This flag tells IPv6 clients to use the router's information to only configure routing and use traditional configuration mechanisms such as DHCPv6 to fetch the IPv6 address.

The Other Stateful Configuration flag (or the *o* flag) tells IPv6 clients to use the router's information to configure routing and build an IPv6 address, but to use a mechanism such as DHCPv6 to get the other information such as the addresses of DNS servers and the suffixes to use when making DNS queries. This is where the distinction between stateless and stateful configuration comes in. Stateless configuration is where an IPv6 client relies wholly on router solicitation and router advertisements to configure IPv6. Stateful configuration is where an IPv6 client relies on a DHCP server or other mechanisms to configure IPv6.

Figure 2 shows an example of an IPv6 address configured from a response to a router solicitation request. Note that the IPv6 and default gateway addresses look very different. This is because Windows uses the link-local address of the router as the default gateway. This is very different from IPv4, where you can manually set the address of the default gateway to its non-link-local address (which in this case is 2001:470:b:a6b::1).

Unfortunately, Windows 8 and Server 2012 are poor IPv6 clients. Windows ignores the addresses of DNS servers and search suffixes

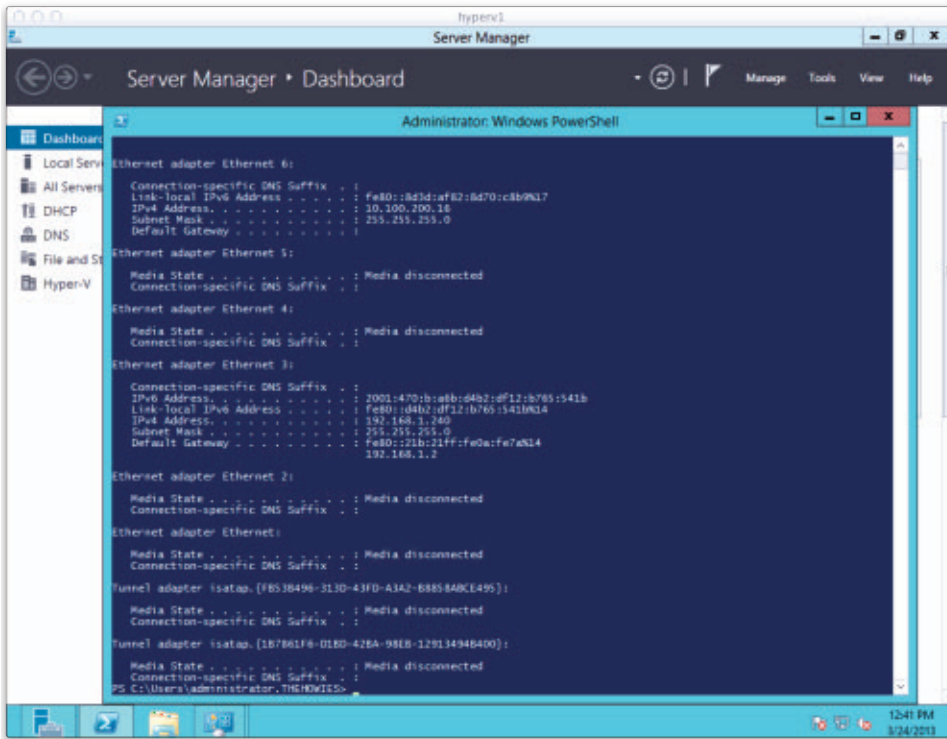


Figure 2
IPv6 Address
Configured from a
Response to a Router
Solicitation Request

provided by IPv6 routers in solicitation responses and advertisements. Even when the *m* flag is set, Windows uses the information provided by a router to build an IPv6 address. And even when told not to use DHCPv6 for other information when the *o* flag is set, it will. In other words, when a router responds to a router solicitation or when a router advertisement is heard, Windows uses the information to build an IPv6 address for the adapter on which the information is heard and still looks for a DHCPv6 server. If a DHCPv6 server is available but doesn't offer IPv6 addresses (i.e., it's configured as a stateless DHCPv6 server set up to serve clients with the *o* flag set and return only DNS server addresses and search suffixes), Windows ignores it. However, if the DHCPv6 server returns an IPv6 address along with DNS server addresses and search suffixes, Windows adds the address to the interface and uses the additional information. That means your Windows

system now has two IPv6 addresses and can use and can be reached on either address. Worse, both addresses will be published in DNS.

Because Windows 8 and Server 2012 always check for a DHCPv6 server, you might wonder why I didn't recommend using DHCPv6.

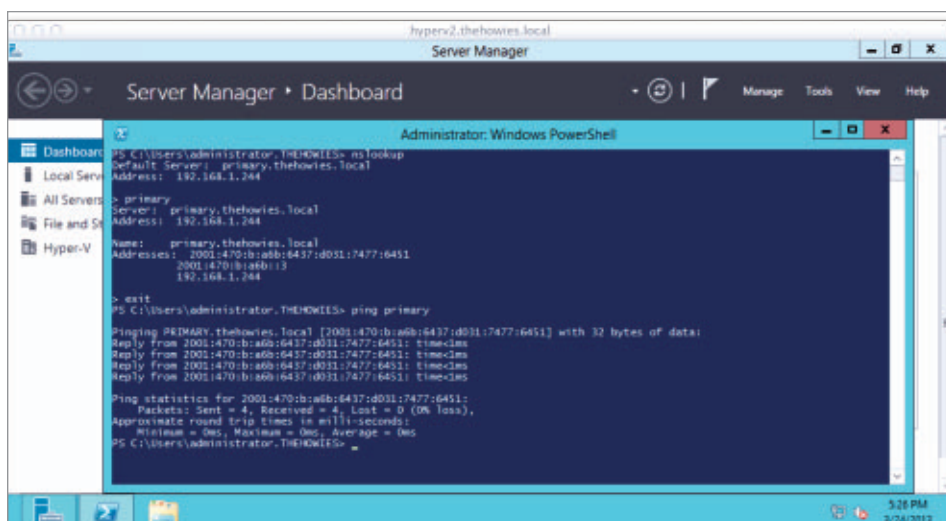
- The IPv6 addresses returned by the DHCPv6 server don't contain enough information by themselves to be usable. They're missing prefix information. Depending on the IPv6 addresses you configure, you might find that Windows assumes the IPv6 prefix is 128-bits, meaning that the host can only communicate with itself.
- In DHCPv6, there's no way to specify the default gateway address. As a result, Windows has to rely on router advertisements to find the IPv6 routers and build a routing table.

My recommendation is that you simply rely on router solicitation and discovery to obtain an auto-configured address and find the default gateway, and use IPv4 to query DNS servers. This setup works well.

Figure 3 shows a Server 2012 system sending ICMPv6 echo requests to a host named Primary, even though Primary has an IPv4 address. Once Server 2012 has an IPv6 address other than its link-local address, it will attempt to use IPv6 by default.

Figure 3

Server 2012 System
Sending ICMPv6 Echo
Requests



Connectivity Testing and Troubleshooting

Testing IPv6 in your network is like testing a sports car in the city. It's necessary, but it's only the first step. You also need to try it out on the information highway.

Figure 4 shows a simple echo request and reply to an IPv6-capable website. As you can see, the Ping command includes the -6 flag, which forces Ping to use IPv6. If all goes well, you should see a reply. If you have a native IPv6 connection to the Internet, the response should be quite speedy. In the example shown in Figure 4, the response time is quite high, because I'm running IPv6 in an IPv4 tunnel with a tunnel broker (i.e., a company that provides IPv6 connectivity). If your echo request fails to elicit a reply, there might be a firewall or other networking device blocking ICMPv6 somewhere between your Windows system and the target.

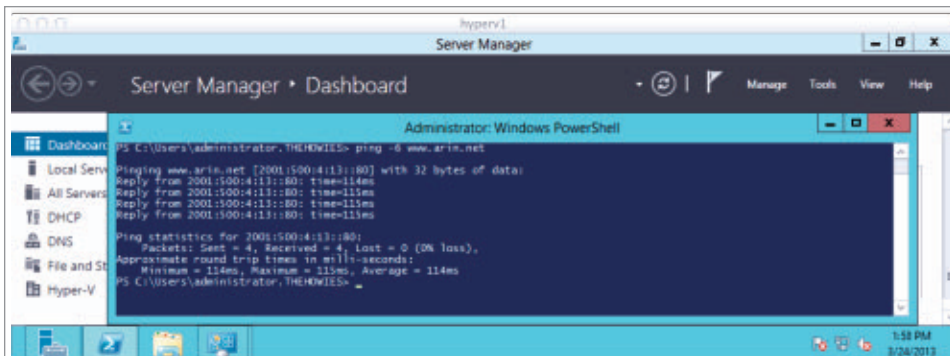
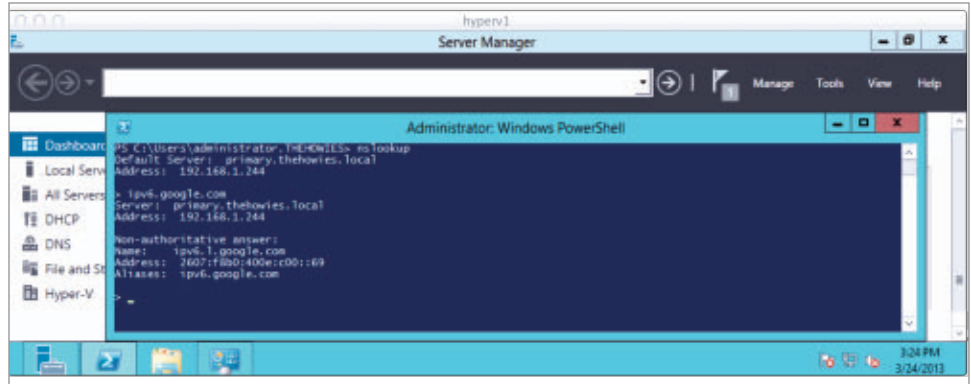


Figure 4
Echo Request
and Reply to an
IPv6-Capable Website

When using firewalls and routers, you need to configure them with rules largely similar to those used for IPv4 networks. Your existing IPv4 rules won't work for the most part. The exception is when the rules are network-layer independent and focus on transport-layer protocols (TCP or UDP) and ports. Whatever you do, no matter how tempted you are, don't configure an IPv6 default rule that allows all traffic to flow between IPv6 interfaces in order to troubleshoot IPv6 connectivity. Cyber criminals, cyber terrorists, and nation states engaged in cyber warfare activities are all proficient in using IPv6.

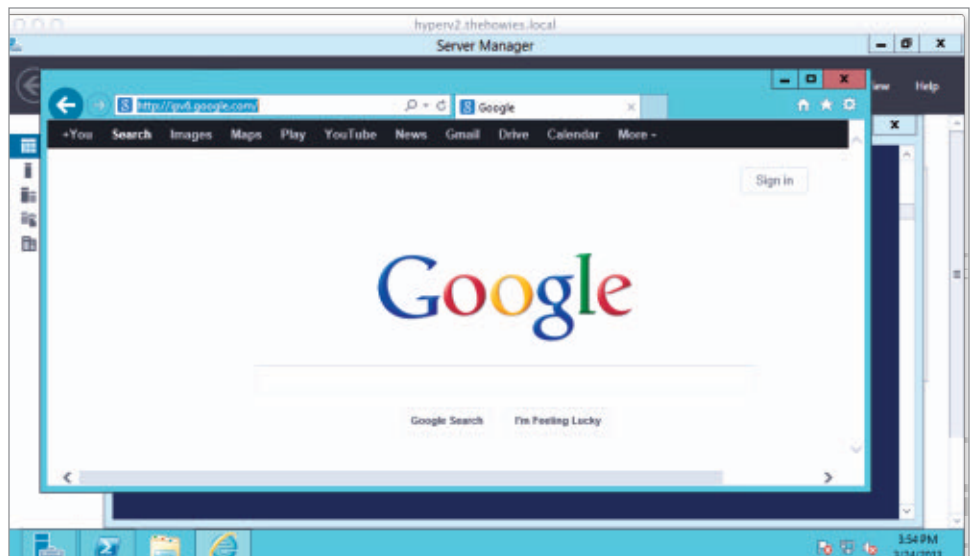
Once you know you have connectivity to the Internet using IPv6, you'll want to test some IPv6-only websites to verify that everything works. Figure 5 shows a DNS lookup for the host `ipv6.google.com`.

Figure 5
DNS Lookup



(In case you're wondering why I used a Google website to test IPv6, Microsoft doesn't offer a website dedicated to IPv6 testing.) As you can see in Figure 5, the lookup came back with only an IPv6 address. Figure 6 shows a browser connected to the `ipv6.google.com` website.

Figure 6
Browser Connected
to an IPv6 Website



Security Concerns

Windows 8 and Server 2012 are particularly adept at obtaining an IPv6 address in a variety of situations and using IPv6 to communicate by default. This can be very problematic in certain environments. A staged migration to IPv6 might have the routers and firewalls configured to support IPv6 and offer router advertisements, but have the *m* and *o* flags configured to prevent clients from using them. Unfortunately, Windows will use the advertisements regardless, and IPv6 communications will begin. Most enterprises processing sensitive data (e.g., Social Security numbers, credit card data) will be using sophisticated system and network monitoring tools, such as intrusion prevention systems (IPSs) and Security Event and Incident Management (SEIM) systems. However, IPv6 support for these types of tools isn't great, and you might find that they're unable to detect suspicious and malicious activity taking place over IPv6. So, before you turn on IPv6 on your Windows networks, make sure that your third-party tools and packages, such as IPSs and SEIM systems, will support it.

Of great concern to many organizations is something called the Advanced Persistent Threat (APT). This term is definitely loaded, but to many people it simply means a very sophisticated attacker who has breached their systems and networks and is able to snoop on data at will. Tools are becoming readily available to deal with APT, but unfortunately they're usually insufficient because they don't account for the use of IPv6 to exfiltrate data from corporate networks. The sheer number of tunnel brokers, Teredo servers, and 6to4 hosts on the Internet makes it infeasible to configure edge defenses to block traffic to all of them.

Of equal to concern to most organizations is the ability to block employees' access to unauthorized websites and cloud services. These organizations often deploy solutions that block the unauthorized sites and services to ensure compliance with statutory and regulatory compliance obligations, prevent accidental infections and data leakage, and increase worker productivity. Such solutions typically



Learning Path

Related Articles

["IPv6 Overview"](#)

["The Inevitability of IPv6, Part 1"](#)

["The Inevitability of IPv6, Part 2"](#)

["IPv6: No Sticks, Just Carrots"](#)

["Managing Your Migration and Transition from IPv4 to IPv6"](#)

["Supporting IPv6 in Your Windows Server 2008 Environment"](#)

["Hands-On IPv6 Lab Setup"](#)

rely on the use of proxy servers or firewalls, and assume that web browsing is taking place over IPv4. These organizations will need to deploy IPv6 gateways and establish similar defenses for IPv6.

If you haven't already done so, I recommend that you regularly inspect network traffic for IPv6 traffic. You should also check your DNS servers for IPv6 addresses that aren't link-local addresses (i.e., addresses that begin with something other than FE80::).

A Challenging Situation for IT

Microsoft has invested a lot of energy into making sure Windows 8 and Server 2012 are able to work in IPv6-ready environments. In fact, the behavior of the Windows IPv6 client might be more about ensuring connectivity than faulty software. From a technical standpoint, this makes Windows one of the best-prepared OSs for IPv6 environments. From a corporate IT standpoint, it creates some challenges. However, with a little bit of planning, IPv6 works great. ■

FAQ

Answers to Your Questions

Q: Can I install the System Center Operations Manager agent in virtual machines in Windows Azure Infrastructure as a Service and monitor them?

A: Yes. A virtual machine (VM) running in Windows Azure Infrastructure as a Service (IaaS) is essentially just a VM running an OS, and you can install the System Center Operations Manager agent within the guest OS and treat it like any other system. As with any other remote location, you need to make sure certain considerations are addressed:

- **Communication.** Can the Operations Manager Management Server communicate with the Operations Manager agents? This means either enabling a site-to-site VPN between your intranet and Azure network or ensuring the correct ports are open between the Operations Manager Management Server and the agents.
- **Authentication.** Are the Azure IaaS VMs part of the domain allowing Kerberos, or will you need to use certificates? Another option would be to install an Operations Manager Gateway in the Azure network, which could use Kerberos to talk to the IaaS VMs—then the gateway would use certificates to talk to the on-premises Operations Manager Management Server.

Remember that Operations Manager also has other types of monitoring for Azure available. This includes more than just agents running inside VMs.

—John Savill



John Savill



Jan De Clercq

Q: Can Microsoft Security Configuration Manager automatically check for security baseline updates on the Microsoft website?

A: Yes, Microsoft Security Configuration Manager (SCM) can automatically check the Microsoft website for updates to the default Microsoft baselines. Default Microsoft baselines are automatically included when you first install SCM on your machine. By default, SCM automatically checks for baseline updates each time you start the tool. You can control SCM's update checking behavior from the Options menu item in SCM's File menu. When you open this menu item, you'll see that the *Check for updates automatically at startup* option is selected by default.

You can also manually check for updates to the default Microsoft baselines. To do so, use the *Check for Updates* option in SCM's File menu. For more information about using baselines in SCM, see [“Adding Settings to Custom Security Baselines in Security Compliance Manager”](#) and [“Comparing Custom and Default Security Baselines in Security Compliance Manager.”](#)

—Jan De Clercq

Q: We want to use SCM to lock down our Windows servers. We duplicated some of the baseline templates that Microsoft provides to create our own custom SCM security baselines. How can we add a specific registry setting to these custom security baselines?

A: To add a specific registry setting to your custom SCM baseline, follow these steps:

1. Start SCM.
2. Go to the Baselines Library pane on the left. In the Custom Baselines section, click the custom baseline to which you want to add the setting.

- Go to the Action pane on the right side. In the Setting area, click Add to display the Add Settings dialog box, which you can see in Figure 1.

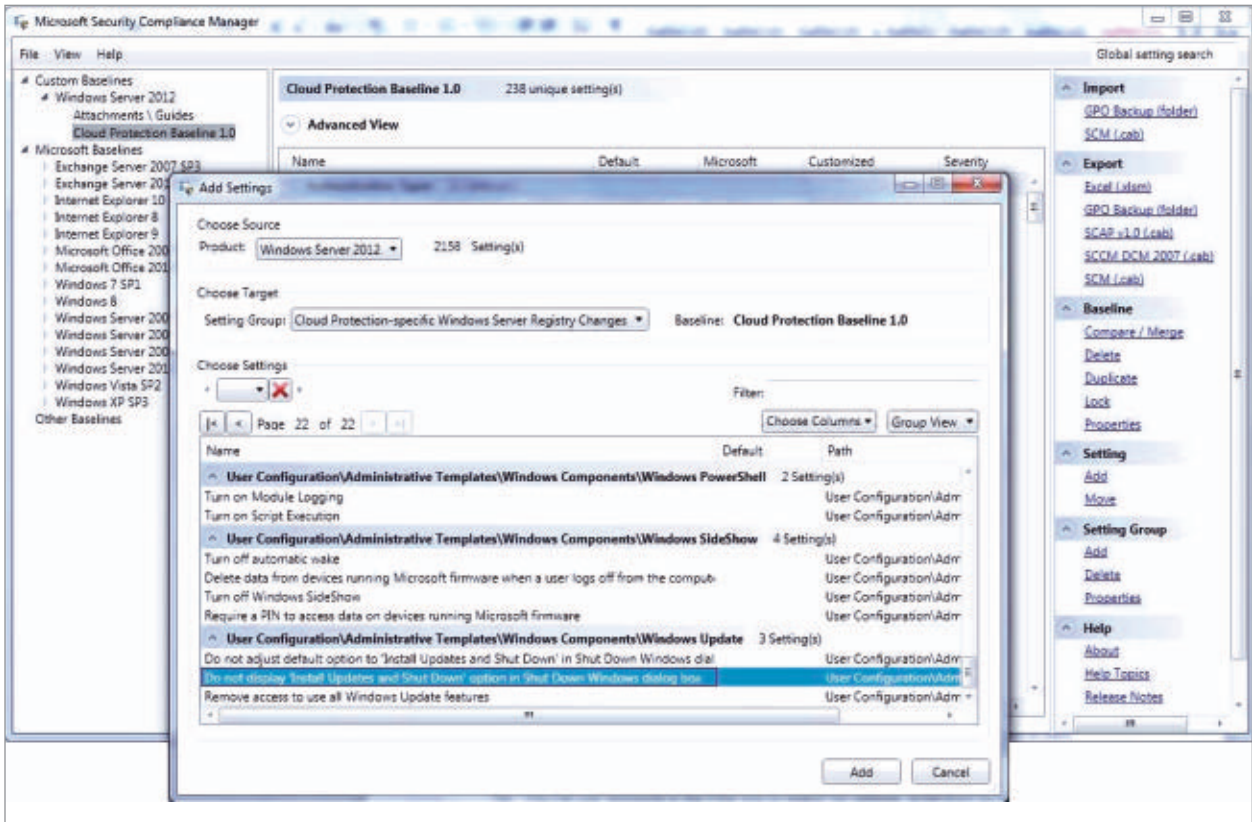


Figure 1

Adding a Setting to a Custom Baseline in SCM

- The settings that SCM displays in the Add Settings dialog box are part of the custom baseline. In this example, the custom baseline is named Cloud Protection Baseline 1.0. It controls the security settings of the Windows servers in my private cloud platform. It's actually a duplicate of the default [Windows Server 2012 File Server](#) baseline, which is why SCM shows *Windows Server 2012* in the Choose Source area.
- In the Choose Settings section, locate the setting that you want to add. Click the setting to select it, then click Add. For this

example, I'm adding a registry key that controls the visibility of the *Install Updates and Shut Down* option in the Windows Shut Down dialog box to the custom baseline.

5. Configure the setting. After you add the setting, SCM takes you back to its standard view, where the new setting appears, as Figure 2 shows. Notice that the new setting isn't configured yet. To configure the *Do not display 'Install Updates and Shut Down' option* setting, simply select the Enabled radio button.

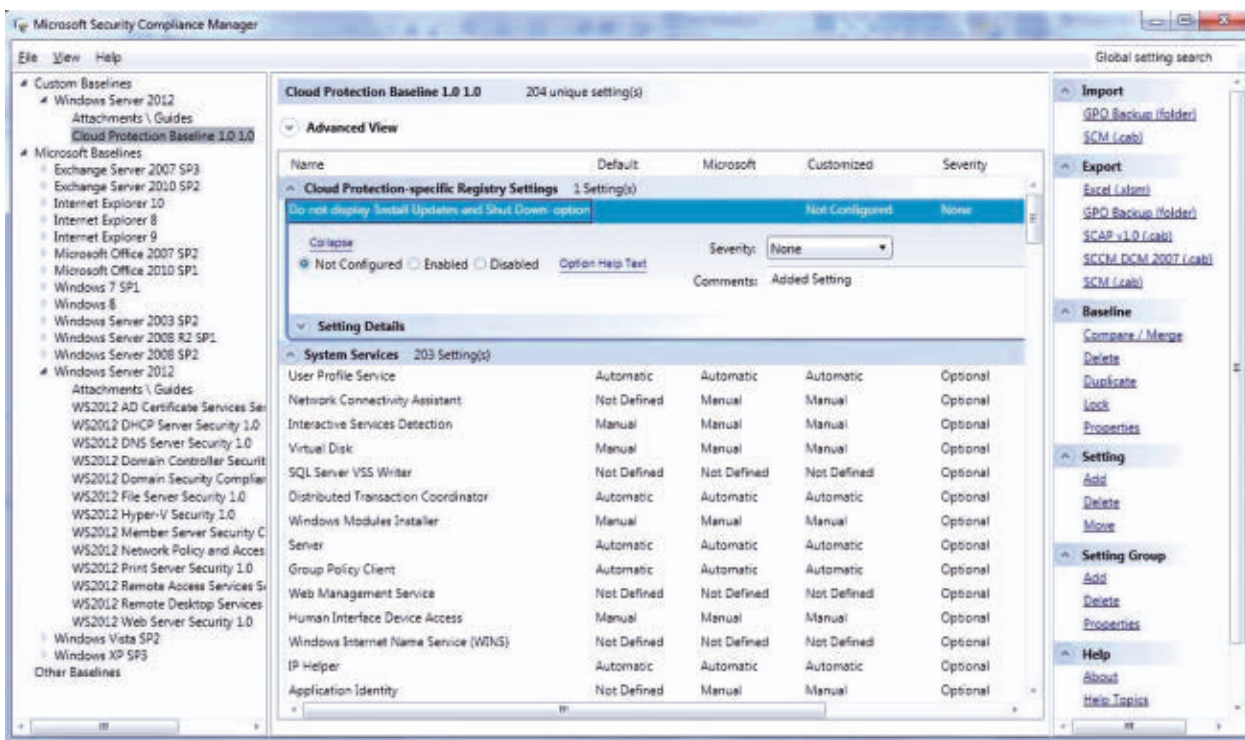


Figure 2

Configuring a Setting for a Custom Baseline in SCM

Note that you can create custom Setting Groups to group the settings you want to control with the SCM baseline. For example, I created a custom Setting Group named *Cloud Protection-specific Windows Server Registry Settings*. As Figure 1 shows, custom Setting Groups appear in the Choose Target area of the Add Settings dialog box. To add a setting to a custom Setting Group, you must expand the Setting

Group drop-down list and select the Setting Group to which you want to add the setting. For more information about using baselines in SCM, check out “[Updating the Default Security Baselines in Security Configuration Manager](#)” and “[Comparing Custom and Default Security Baselines in Security Compliance Manager](#).”

—Jan De Clercq

Q: What’s the goal of the primary computer feature that Microsoft introduced in Windows Server 2012 Active Directory? How can I leverage this feature to better protect our corporate data?

A: The primary computer feature allows Active Directory (AD) administrators to label AD computer objects as the primary computers of certain domain users. AD administrators can use this feature to specify the computers on which users’ roaming profiles can be downloaded and specify the computers on which users can get access to their redirected folders. When users log on to computers that haven’t been labeled as primary computers, they’ll get a local profile, and they won’t get access to their redirected folders.

In this age of the consumerization of IT and trends such as bring your own device (BYOD), using the primary computer feature is a powerful way to associate or dissociate user data and settings with particular computers or devices. Designating primary computers reduces the security and privacy risks of downloading or leaving personal and corporate data on personal or public computers on which users have logged on.

The primary computer feature is based on a set of new Group Policy Object (GPO) settings and an AD schema extension. When a user logs on to a [Windows 8](#) or Server 2012 machine, the logon logic will check the status of two GPO-controlled settings to determine whether the msDS-Primary-Computer attribute that’s linked to the AD user account object of the user who is logging on should

influence the decision to roam the user's profile or apply folder redirection. The two GPO settings are:

- *Download roaming profiles on primary computers only*, which is located in the \User Configuration\Policies\Administrative Templates\System\User Profiles GPO container
- *Redirect folders on primary computers only*, which is located in the \User Configuration\Policies\Administrative Templates\System\Folder Redirection GPO container

You can use the Active Directory Administrative Center or Windows PowerShell cmdlets to populate an AD user object's msDS-Primary-Computer attribute with the distinguished names (DNs) of computer accounts that should be marked as a user's primary computers.

The support for the primary computer feature requires that your AD schema is upgraded to Server 2012. It can only be leveraged on domain-joined Server 2012 and Windows 8 machines. For more details on how to set this up, I recommend that you read the Microsoft Storage Team blog post "[Configuring Primary Computers for Folder Redirection and Roaming Profiles in Windows Server '8' Beta.](#)"

—Jan De Clercq

Q: What are the scalability limits for System Center 2012 Virtual Machine Manager SP1?

A: System Center 2012 Virtual Machine Manager (SCVMM) SP1 increased the scalability for management nearly three-fold over the previous version. It provides the following for each VMM management server: 1,000 hosts; 25,000 virtual machines (VMs); and 64-node cluster support. ■

—John Savill

Product News for IT Pros

Dell Enhances Management and Monitoring of Microsoft Environments



Dell announced a series of enhancements that are designed to help organizations optimize the migration, management, and monitoring of Microsoft environments. New releases include the latest release of Spotlight for SQL Server Enterprise, Spotlight Project Lucy (a new cloud-based set of productivity tools that enable SQL Server users to obtain a free system health check), enhancements to Dell KACE K2000 deployment appliances, and a new release of MessageStats Business Insights. With the release of Spotlight for SQL Server Enterprise 9.5, DBAs can now monitor the health and performance of their SQL Server environments anytime, anywhere, and from any device. The new release features a Windows 8 mobile app, which allows customers to view their performance heat map and receive alerts directly from their Windows phones and tablets. For more information, visit the [Dell website](#).

KineticD Introduces SharePoint Backup Protection



KineticD announced that its KineticCloud Backup for Servers has been enhanced to protect all data associated with Microsoft SharePoint, including content databases, service applications, and search data, leveraging Microsoft's native APIs. Small-to-midsized businesses (SMBs) and channel partners can set up granular backups for even greater backup and recovery efficiency. KineticD sends incremental data changes and only transfers blocks that haven't been stored in the vault, eliminating unnecessary data at the source and providing much faster backups that consume less storage, bandwidth, and time. Also, backups don't require installing individual agents or plug-ins,

providing greater scalability and efficiency. Backups can be scheduled for off-peak times, or at specific intervals during the day for better workflow. The product also offers hybrid cloud protection, giving you access to offline (local) and online (remote) versions of files to provide ongoing availability to digital assets at all times. For more information, check out the [KineticD website](#).



Devolutions Updates Remote Desktop Manager

Devolutions released Remote Desktop Manager 8.3, the new version of its all-in-one management platform for IT teams. The product lets you centralize all your remote connections, passwords, and credentials in a usable platform. Remote Desktop Manager 8.3 features more than 80 improvements and bug fixes, including support for new data source types (SFTP, FTPS, MariaDB), an improved UI (e.g., Google Maps and cell phone field in the Contacts list, simplified application password management), Secret Server Windows authentication support, and an improved import connection framework (support for Royal TSX and Terminals 3.0 file types). For more information, visit the [Devolutions website](#).



KEMP Provides Smart Load Balancing

KEMP Technologies builds server load-balancing application delivery controllers (ADCs) that are becoming increasingly popular with small-to-midsized businesses (SMBs) and now—with its more enterprise-focused solutions—larger companies as well. April saw the release of the LoadMaster R320 purpose-built load balancer, based on the Dell PowerEdge R320 server platform. The partnership with Dell enabled key optimizations through a couple of Dell capabilities, including the Integrated Dell Remote Access Controller (iDRAC) and the Integrated Lifecycle Controller. These features let users manage and monitor the Loadmaster and Dell servers through a common enterprise platform. Customers can easily track hardware problems and perform remote management. Now, to further entice the enterprise crowd—and to fill

a gap left behind after Microsoft discontinued Forefront Threat Management Gateway (TMG)—KEMP has added new security features to its load balancers with its new Edge Security Pack for the KEMP Load-Master line. The Edge Security Pack (ESP) helps companies protect business-critical, web-facing applications from unauthorized access. For more information, check out the [KEMP Technologies website](#).

Fujitsu Launches Enterprise-Class Premium Notebooks

Fujitsu announced its new LIFEBOOK E Series family of business notebook computers, designed to optimize enterprise mobility. The stylish and lightweight premium notebook PCs are ideal for demanding business users who are constantly on the go, require reliable devices, and want to save on costs—but not compromise on features. All three of the new LIFEBOOK models offer full business functionality and convenient functions that are expected in today's premium business notebooks. The LIFEBOOK E Series is available as 13", 14", and 15.6" models, and includes a roomy touchpad with integrated buttons, an optional backlit keyboard, and the Fujitsu signature modular bay that is capable of hosting a second battery, a second hard disk drive, an optical drive, or a weight-saver module to give users the lightest possible traveling weight. For more information, see the [Fujitsu website](#).



CommVault Simpana 10 Offers Massively Scalable, Open Software Platform

With the latest release of its data management software platform, Simpana 10, CommVault is enabling enterprises to take an exponential leap forward in protecting and managing their data. Simpana 10 extends CommVault's data protection and archiving leadership to deliver secure, self-service access from mobile devices, speed the adoption of cloud computing, and extract value from Big Data. For the first time, employees across the enterprise—not just IT managers—can easily repurpose data under management and quickly search, access, and create information to enable better decision making



and collaboration. New innovations include enhanced IntelliSnap snapshot management (providing instant, automated recovery of applications and virtual servers); Simpana OnePass with Exchange (making long-term email retention affordable and practical by converging backup, archiving, and reporting into a single process); and tighter integration with Microsoft Hyper-V, VMware vSphere 5.1, and vCloud Director 5.1 (helping enterprises achieve cloud-scale through automated discovery, protection, and recovery in virtualized environments). For more information, see the [CommVault website](#).



Nasuni Provides a New View of Enterprise Data

Nasuni announced the Nasuni Management Console (NMC), an easy-to-use virtual appliance that provides a central command center from which to manage all of an organization's data regardless of its physical location, empowering true IT storage agility and unparalleled data visibility. With the NMC, IT can manage data access, protection, security, and storage capacity from a single console, without the need to do so through individual hardware components. Until now, enterprises couldn't manage storage effectively outside of the primary data center. Managing data storage at remote and branch offices requires logging on to a wide array of different machines or putting someone on a plane. Data protection, in particular, is a nightmare for the distributed enterprise. The NMC is an integral part of Nasuni's Storage-Infrastructure-as-a-Service (SIaaS), which gives enterprise organizations a secure, all-in-one data storage solution that provides local performance for users, simplified and centralized management for IT, and an easily scalable storage solution for the enterprise that can save organizations as much as 60 percent over traditional hardware solutions. For more information, visit the [Nasuni website](#). ■

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
forums.windowsitpro.com

News

Check out the current news and information about Microsoft Windows technologies.
www.winsupersite.com

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- Dev Pro UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- SharePoint Pro UPDATE
- SQL Server Pro UPDATE
- Windows IT Pro UPDATE
- WinInfo Daily UPDATE

RELATED PRODUCTS

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
windowsitpro.com/vip-premium-membership

SQL Server Pro

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

Dev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
www.sharepointpromag.com

Advertiser Directory

1&1 Internet	1
IT/Dev Connections	57
Windows IT Pro	2, 49

Vendor Directory

CommVault	97, 98
Concur Technologies	21
Dell	95, 96

Devolutions	96
Fujitsu	97
Google	19, 21, 22
KEMP Technologies	96, 97
KineticD	95, 96
Nasuni	98
Ping Identity	21
Twitter	21, 22